arXiv:1611.01647v1 [cs.DS] 5 Nov 2016

UNIFORM SAMPLING THROUGH THE LOVÁSZ LOCAL LEMMA

HENG GUO, MARK JERRUM, AND JINGCHENG LIU

Abstract. We propose a new algorithmic framework, called "partial rejection sampling", to draw samples exactly from a product distribution, conditioned on none of a number of bad events occurring. Our framework builds (perhaps surprising) new connections between the variable framework of the Lovász Local Lemma and some classical sampling algorithms such as the "cycle-popping" algorithm for rooted spanning trees by Wilson [28]. The design of our algorithms is significantly different from existing sampling techniques. Among other applications, we discover new algorithms to sample satisfying assignments of k-CNF formulas with bounded variable occurrences.

1. Introduction

The Lovász Local Lemma [9] is a classical gem in combinatorics that guarantees the existence of a perfect object that avoids all events deemed to be "bad". The original proof is non-constructive but there has been great progress in the algorithmic aspects of the local lemma. After a long line of research [3, 2, 19, 8, 26], the celebrated result by Moser and Tardos [20] gives efficient algorithms to find such a perfect object under conditions that match the Lovász Local Lemma in the so-called variable framework. However, it is natural to ask whether, under the same condition, we can also sample a perfect object uniformly at random instead of merely finding one.

Roughly speaking, the resampling algorithm by Moser and Tardos [20] works as follows. We initialize all variables randomly. If bad events occur, then we randomly choose a bad event and resample all the involved variables. Unfortunately, it is not hard to see that this algorithm can produce biased samples. This seems inevitable. As Bezáková et al. showed [4], sampling can be **NP**-hard even under conditions that are stronger than that of the local lemma. On the one hand, the symmetric Lovász Local Lemma only requires $ep\Delta \leq 1$, where p is the probability of bad events and Δ is the maximum degree of the dependency graph. On the other hand, translating the result of [4] to this setting, one sees that as soon as $p\Delta^2 \geq C$ for some constant C, then even approximately sampling perfect objects in the variable framework becomes **NP**-hard.

The starting point of our work is a new condition (see Condition 5) under which we show that the output of the Moser-Tardos algorithm is in fact uniform (see Theorem 8). Intuitively, the condition requires any two dependent bad events to be disjoint. Indeed, instances satisfying this condition are called "extremal" in the study of Lovász Local Lemma. For these extremal instances, we can in fact resample in a parallel fashion, since the occurring bad events form an independent set in the dependency graph. We call this algorithm "partial rejection sampling", in the sense that it is like rejection sampling, but only resamples the bad events.

Our result puts some classical sampling algorithms under a unified framework, including the "cycle-popping" algorithm by Wilson [28] for sampling rooted spanning trees, and the "sink-popping" algorithm by Cohn, Pemantle, and Propp [7] for sampling sink-free orientations of an undirected graph. Indeed, Cohn et al. [7] coined the term "partial rejection sampling" and asked for a general theory, and we believe that extremal instances under the variable framework is a satisfactory answer. With our techniques, we are able to give a new algorithm to sample solutions for a special class of k-CNF formulas, under conditions matching the Lovász Local

HG and MJ are supported by the EPSRC grant EP/N004221/1. JL is supported by NSF grant CCF-1420934. This work was done (in part) while the authors were visiting the Simons Institute for the Theory of Computing.

Lemma (see Corollary 20), which is an **NP**-hard task for general k-CNF formulas. Furthermore, we provide *explicit* formulas for the expected running time of these algorithms (see Theorem 13), which matches the running time upper bound given by Kolipaka and Szegedy [16] under Shearer's condition [24].

The next natural question is thus whether we can go beyond extremal instances. Indeed, we also give a generalized uniform sampler that applies to *any* problem under the variable framework. The main idea is that, instead of only resampling occurring bad events, we resample a larger set of events so that the choices made do not block any perfect assignments in the end, in order to make sure of uniformity in the final output.

As a simple example, we describe how our algorithm samples independent sets. The algorithm starts by choosing each vertex with probability $1/2$ independently. At each subsequent round, in the induced subgraph on the currently chosen vertices, the algorithm finds all the connected components of size ≥ 2. It marks these vertices and their boundaries (which are unoccupied) as "to be resampled". Then it resamples all marked vertices, and repeats this process until there is no edge with both endpoints occupied. What seems surprising is that this simple process does yield a uniformly random independent set when it stops. Indeed, as we will show in Theorem 35, this simple process is an exact sampler for *weighted independent sets* (also known as the *hard-core model* in statistical physics). In addition, it runs in expected linear time under a condition that matches, up to a constant factor, the *uniqueness threshold* of the model (beyond which the problem of approximate sampling becomes **NP**-hard).

In the more general setting, we will mark the set of events to be resampled, denoted by Res, iteratively. We start from the set of occurring bad events. Then we mark neighbouring events of the current set Res, until there is no event A on the boundary of Res such that the current assignment, projected on the common variables of A and Res, can be extended so that A may happen. In the worst case, we will resample all events (there is no event in the boundary at all). In that scenario the algorithm is the same as a naive rejection sampling, but typically we resample many fewer variables in every step. We show that this is a uniform sampler on assignments that avoid all bad events once it stops (see Theorem 25).

One interesting feature of our algorithm is that, unlike Markov chain based algorithms, ours does not require the solution space (or any augmented space) to be connected. Moreover, our sampler is exact; that is, when the algorithm halts, the final distribution is precisely from the desired distribution. Prior to our work, most exact sampling algorithms are obtained by coupling from the past [21].

We also give sufficient conditions that guarantee a linear expected running time of our algorithm in the general setting (see Theorem 26). The first condition is that $p\Delta^2$ is bounded above by a constant. This is optimal up to constants in observance of the **NP**-hardness result in [4]. Unfortunately, this condition alone does not make the algorithm efficient. In addition, we also need to bound the expansion from bad events to resampling events, which leads to an extra condition on intersections of bad events.

To illustrate the result, we apply our algorithm to sample satisfying assignments of k-CNF formulas. For k-CNF formulas, the extra condition naturally leads to a lower bound on the number of shared variables between two dependent clauses. We say that a k-CNF formula has intersection s if any two *dependent* clauses share at least s variables. Let n be the number of variables. We (informally) summarize our results on k-CNF formulas as follows (see Corollary 31 and Theorem 33):

- If $d \leq \frac{1}{\sqrt{6ek}} \cdot 2^{k/2} + 1$ and $s \geq \log(d-1) + \log k + \log(3e)$, then our generalized partial rejection resampling algorithm outputs a uniformly random solution to a k-CNF formula with degree d and intersection s in expected running time $O(n)$.
- If $d \geq 4 \cdot 2^{k/2}$ (for an even k), then even if $s = k/2$, it is **NP**-hard even to *approximately* sample a solution to a k-CNF formula with degree d and intersection s.

As shown in the hardness result, the intersection bound does not render the problem trivial.

Prior to our work, the best known result of sampling/counting satisfying assignments of k-CNF formulas is from [4] which requires the formula to be monotone and $d \leq k$ to be

large enough (see also [5, 17]).[1] Although our result requires an additional lower bound on intersections, not only does it improve the dependency of k and d *exponentially*, but also achieves a *matching* constant $1/2$ in the exponent. Furthermore the samples produced are exactly uniform. Thus, if the extra condition on intersections can be removed, we will have a sharp phase transition at around $d = \tilde{O}(2^{k/2})$ in the computational complexity of sampling solutions to k-CNF formulas with bounded variable occurrences. Such a sharp transition has been recently established for, e.g., sampling configurations in the hard-core model [27, 25, 10].

While the study of algorithmic Lovász Local Lemma has progressed beyond the variable framework [12, 1, 13], we restrict our focus to the variable framework in this work. It is also an interesting future direction to investigate and extend our techniques of uniform sampling beyond the variable framework. For example, one may want to sample a permutation that avoids certain bad events. The classical sampling problem of perfect matchings in a bipartite graph can be characterized in this way.

2. Partial Rejection Sampling

We first describe the "variable" framework. Let $\{X_1, \ldots, X_n\}$ be a set of random variables. Each X_i can have its own distribution and range D_i. Let $\{A_1, \ldots, A_m\}$ be a set of "bad" events that depend on X_i's. For example, for a constraint satisfaction problem (CSP) with variables X_i $(1 \le i \le n)$ and constraints C_j $(1 \le j \le m)$, each A_j is the set of unsatisfying assignments of C_j for $1 \le j \le m$. Let $\mathsf{var}(A_i)$ be the (index) set of variables that A_i depends on.

The dependency graph $G = (V, E)$ has m vertices, identified with the integers $\{1, 2, \ldots, m\}$, corresponding to the events A_i, and (i, j) is an edge if A_i and A_j depend on one or more common variables. In other words, $(i, j) \in E$ if $\mathsf{var}(A_i) \cap \mathsf{var}(A_j) \neq \emptyset$. We write $A_i \sim A_j$ if the vertices i and j are adjacent in G. The asymmetric Lovász Local Lemma [9] states the following.

Theorem 1. *If there exist non-negative real numbers $x_i < 1$ $(1 \le i \le m)$ such that $\forall i$,*

$$\Pr(A_i) \le x_i \prod_{(i,j)\in E} (1 - x_j), \tag{1}$$

then $\Pr\left(\bigwedge_{i=1}^{m} \overline{A_i}\right) \ge \prod_{i=1}^{m}(1 - x_i) > 0.$

Theorem 1 has a condition that is easy to verify, but not necessarily optimal. Shearer [24] gives the optimal condition for the local lemma to hold for a fixed dependency graph G. To state Shearer's condition, we will need the following definitions. Let $p_i := \Pr(A_i)$ for all $1 \le i \le m$. Let $\mathcal{I}$ be the collection of independent sets of G. Define the following quantity:

$$q_I(\mathbf{p}) := \sum_{J\in\mathcal{I},\, I\subseteq J} (-1)^{|J|-|I|} \prod_{i\in J} p_i,$$

where $\mathbf{p} = (p_1, \ldots, p_m)$. When there is no confusion we also simply write q_I instead of the more cumbersome $q_I(\mathbf{p})$. Note that if $I \notin \mathcal{I}$, $q_I = 0$.

Theorem 2 (Shearer [24])**.** *If $q_I \ge 0$ for all $I \subseteq V$, then* $\Pr\left(\bigwedge_{i=1}^{m} \overline{A_i}\right) \ge q_\emptyset$.

In particular, if the condition holds with $q_\emptyset > 0$, then $\Pr\left(\bigwedge_{i=1}^{m} \overline{A_i}\right) > 0$.

Neither Theorem 1 nor Theorem 2 yields an efficient algorithm to find the assignment avoiding all bad events, since they only guarantee an exponentially small probability. There has been a long line of research devoting to an algorithmic version of LLL, culminating in Moser and

[1] Shortly before posting our paper, Hermon, Sly, and Zhang [14] posted a result improving this bound to a tight $d \le c2^{k/2}$ for monotone k-CNF formulas. At about the same time, Moitra [18] posted an algorithm for general k-CNF when $d \le 2^{k/20}$. We note that both techniques are quite different from ours and the results are not directly comparable. In particular, both of these samplers are approximate while ours is exact, despite the differences in the range of parameters and whether the formula is monotone or not. Moreover, our algorithm works best when clauses share many variables, which is completely the opposite to the Markov chain algorithm of [14].

Tardos [20] with essentially the same condition as in Theorem 1. The Resample algorithm of Moser and Tardos is very simple, described in Algorithm 1.

Algorithm 1 The Resample algorithm

(1) Draw independent samples of all variables $X_1, \dots, X_n$ from their respective distributions.

(2) While at least one A_i holds, uniformly at random pick one of such A_i and resample all variables in $\mathsf{var}(A_i)$.

(3) Output the current assignment.

In [20], Moser and Tardos showed that Algorithm 1 finds a good assignment very efficiently.

Theorem 3 (Moser and Tardos [20])**.** *Under the condition of Theorem 1, the expected number of resampling steps in Algorithm 1 is at most* $\sum_{i=1}^{m} \frac{x_i}{1-x_i}$.

Unfortunately, the final output of Algorithm 1 is not distributed as we would like, namely as a product distribution conditioned on avoiding all bad events.

In addition, Kolipaka and Szegedy [16] showed that up to Shearer's condition, Algorithm 1 is efficient. To simplify the notation, let $q_i := q_{\{i\}}$ for $1 \le i \le m$.

Theorem 4 (Kolipaka and Szegedy [16])**.** *If* $q_I \ge 0$ *for all* $I \in \mathcal{I}$ *and* $q_\emptyset > 0$*, then the expected number of resampling steps in Algorithm 1 is at most* $\sum_{i=1}^{m} \frac{q_i}{q_\emptyset}$.

On the other hand, Wilson's cycle-popping algorithm [28] is very similar to the Resample algorithm but it outputs a uniformly random rooted spanning tree. Another similar algorithm is the sink-popping algorithm by Cohn, Pemantle, and Propp [7] to generate a sink-free orientation uniformly at random. Upon close examination of these two algorithms, we found a common feature of both problems.

Condition 5. *If* $(i, j) \in E$ *(or equivalently* $A_i \sim A_j$*), then* $\Pr(A_i \wedge A_j) = 0$*; namely the two events* A_i *and* A_j *are disjoint if they are dependent.*

In other words, any two events A_i and A_j are either independent or disjoint. These instances have been noticed in the study of Lovasz Local Lemma. They are the ones that minimize $\Pr\left(\bigwedge_{i=1}^{m} \overline{A_i}\right)$ given Shearer's condition (namely $\Pr\left(\bigwedge_{i=1}^{m} \overline{A_i}\right) = q_\emptyset$). Instances satisfying Condition 5 have been named *extremal* [16].

We will show that, given Condition 5, the final output of the Resample algorithm is a sample from a conditional product distribution (Theorem 8). Moreover, we will show that under Condition 5, the running time upper bound $\sum_{i=1}^{m} \frac{q_i}{q_\emptyset}$ given by Kolipaka and Szegedy (Theorem 4) is indeed the exact expected running time. See Theorem 13.

In fact, when Condition 5 holds, at each step of Algorithm 1, the occurring events form an independent set of the dependency graph G. Think of the execution of Algorithm 1 as going in rounds. In each round we find the set I of bad events that occur. Due to Condition 5, $\mathsf{var}(A_i) \cap \mathsf{var}(A_j) = \emptyset$ for any $i, j \in I$, i.e., I is an independent set in the dependency graph. Therefore, we can resample all variables involved in the occurring bad events without interfering with each other. This motivates Algorithm 2.

We call Algorithm 2 the Partial Rejection Sampling algorithm. This name was coined by Cohn, Pemantle, and Propp [7]. Indeed, they ask as an open problem how to generalize their sink-popping algorithm and Wilson's cycle popping algorithm. We answer this question under the variable framework. Partial Rejection Sampling differs from the normal rejection sampling algorithm by only resampling "bad" events.

In fact, Algorithm 2 can be viewed as a parallel version of Algorithm 1. Suppose each event is assigned to a processor, which determines whether the event holds by looking at the variables associated with the event. If the event holds then all associated variables are resampled. No conflict will be created due to Condition 5.

Algorithm 2 PARTIAL REJECTION SAMPLING

(1) Draw independent samples of all variables $X_1, \dots, X_n$ from their respective distributions.

(2) While at least one bad event holds, find the independent set I of occurring A_i's. Independently resample all variables in $\bigcup_{i \in I} \mathsf{var}(A_i)$.

(3) Output the current assignment.

TABLE 1. A resampling table with 4 variables

X_1	$X_{1,1}$	$X_{1,2}$	$X_{1,3}$	$\dots$
X_2	$X_{2,1}$	$X_{2,2}$	$X_{2,3}$	$\dots$
X_3	$X_{3,1}$	$X_{3,2}$	$X_{3,3}$	$\dots$
X_4	$X_{4,1}$	$X_{4,2}$	$X_{4,3}$	$\dots$

In the following analysis, we will use the resampling table idea, which has appeared in both the analysis of Moser and Tardos [20] and Wilson [28]. Note that we only use this idea to analyze the algorithm rather than to really create the table in the execution. Associate each variable X_i with an infinite stack of random values $\{X_{i,1}, X_{i,2}, \dots\}$. This forms the *resampling table* where each row represents a variable and there are infinitely many columns, as shown in Table 1. In the execution of the algorithm, when a variable needs to be resampled, the algorithm draws the top value from the stack, or equivalently move from the current entry in the resampling table to its right.

Let t be a positive integer to denote the round of Algorithm 2. Let i_t be the index of the variable X_i in the resampling table at round t. In other words, at the t-th round, X_i takes value X_{i,i_t}. Thus, the set $\sigma_t = \{X_{i,i_t} \mid 1 \le i \le m\}$ is the current assignment of variables at round t. This σ_t determines which events happen. Call the set of occurring events, viewed as a subset of the vertex set of the dependency graph, I_t. (For convenience, we shall sometimes identify the event A_i with its index i; thus, we shall refer to the "events in S" rather than the "events indexed by S".) As explained above, I_t is an independent set of G due to Condition 5. Then variables that are involved in any of the events in I_t are resampled. In other words,

$$i_{t+1} = \begin{cases} i_t + 1 & \text{if } \exists j, A_j \in I_t \text{ and } X_i \in \mathsf{var}(A_j); \\ i_t & \text{otherwise.} \end{cases}$$

Therefore, any event that happens in round $t+1$, must share at least one variable with some event in I_t (possibly itself). In other words, $I_{t+1} \subseteq \Gamma^+(I_t)$ where $\Gamma^+(\cdot)$ denotes the set of all neighbours of I unioned with I itself. This inspires the following definition.

Definition 6. *A series $\mathcal{S} = S_1, S_2, \dots, S_\ell$ of independent sets in G is called an* independent set sequence *if one of the following holds:*

(1) $S_i \neq \emptyset$ for all $1 \le i \le \ell$ and for every $1 \le i \le \ell - 1$, $S_{i+1} \subseteq \Gamma^+(S_i)$;
(2) $\ell = 0$ and $\mathcal{S}$ is empty.

Let M be a resampling table. Suppose running Algorithm 2 on M does not terminate up to some integer $\ell \ge 1$ rounds. Define the *log* of running Algorithm 2 on M up to round ℓ as the sequence of independent sets $I_1, I_2, \dots, I_\ell$ created by this run. Thus, for any M and $\ell \ge 1$, the log $I_1, I_2, \dots, I_\ell$ must be an independent set sequence.

Lemma 7. *Suppose Condition 5 holds. Given any log $\mathcal{S} = S_1, S_2, \dots, S_\ell$ of length $\ell \ge 1$, $\sigma_{\ell+1}$ is a random sample from the product distribution conditioned on none of the events A_i occurring, where $i \notin \Gamma^+(S_\ell)$.*

Proof. The set of occurring events at round ℓ is S_ℓ. Hence $\sigma_{\ell+1}$ does not make any of the A_i's happen where $i \notin \Gamma^+(S_\ell)$. Call an assignment σ *valid* if none of A_i's happen where $i \notin \Gamma^+(S_\ell)$. To show that $\sigma_{\ell+1}$ has the desired conditional product distribution, we will show

that the probabilities of getting any two valid assignments σ and σ' are proportional to their probabilities of occurrence in the product distribution.

Let M be the resampling table so that the log of the algorithm is $\mathcal{S}$ up to round $\ell \geq 1$, and $\sigma_{\ell+1} = \sigma$. Indeed, since we only care about events up to round $\ell+1$, we may truncate the table so that $M = \{X_{i,j} \mid 1 \leq i \leq n,\ \ 1 \leq j \leq i_{\ell+1}\}$. Let $M' = \{X'_{i,j} \mid 1 \leq i \leq n,\ \ 1 \leq j \leq i_{\ell+1}\}$ be another table where $X'_{i,j} = X_{i,j}$ if $j < i_{\ell+1}$ for any $i \in [n]$, but $\sigma_{\ell+1} = \sigma'$. In other words, we only change the values in the final round ($X_{i,i_{\ell+1}}$), and only to another valid assignment.

The lemma follows if the algorithm running on M' generates the same log $\mathcal{S}$. Since if this is the case, then conditioned on the log $\mathcal{S}$, every possible table M is one-to-one corresponding to another table M' where $\sigma_{\ell+1} = \sigma'$. Hence the probability of getting σ is proportional to its weight in the product distribution.

Suppose otherwise and the logs obtained by running the algorithm on M and M' differ. Let $t_0 \leq \ell$ be the first round where resampling changes. Since $X'_{i,j} = X_{i,j}$ if $j < i_{\ell+1}$ for any $i \in [m]$, any event A that occurs in $\mathcal{S}$ should still occur when running on M'. (If this is not the case, then A must depend on values at the final round of the resampling table. When the variables of A are resampled, the algorithm will attempt to access values beyond the table, a contradiction.) There must be an occurring event, say A, that happens at t_0 on M' but not on M. Moreover, there must be a set of variables in $\mathsf{var}(A)$ that have values $(X_{i,i_{\ell+1}})$, as otherwise the two runs should be identical. Let us call this set of variables Y. Since resampling does not change before t_0, in the M' run, the assignment of variables in Y must be $(X'_{i,i_{\ell+1}})$ at time t_0.

We claim that $Y = \mathsf{var}(A)$. If the claim does not hold, then $Z := \mathsf{var}(A) \setminus Y \neq \emptyset$. Any variable in Z has not reached final round, and must be resampled in the M run. Let $X_j \in Z$ be the first such variable being resampled at or after round t_0 in the M run. (The choice of X_j may not be unique, and we just choose an arbitrary one.) Recall that $Y \neq \emptyset$, A can no longer happen, thus there must be $A' \neq A$ causing such a resampling of X_j. Then $\mathsf{var}(A) \cap \mathsf{var}(A') \neq \emptyset$. Consider any variable $X_k \in \mathsf{var}(A) \cap \mathsf{var}(A')$. It is resampled at or after time t_0 in the M run due to A'. Hence $X_k \in Z$ for any such k. Moreover, in the M run, until A' happens, X_k has not been resampled since time t_0, because A' is the first resampling event at or after time t_0 that involves variables in Z. On the other hand, in the M' run, X_k's value causes A to happen at time t_0. Hence, there exists an assignment on variables in $\mathsf{var}(A) \cap \mathsf{var}(A')$ such that both A and A' happen. Clearly this assignment can be extended to a full assignment so that both A and A' happen. However, $A \sim A'$ as they share the variable X_j. Due to Condition 5, $A \cap A' = \emptyset$. Contradiction! Therefore the claim holds.

We argue that the remaining case, $Y = \mathsf{var}(A)$, is also not possible. Since A occurs in the M' run, we know, by the definition of σ', that $A \in \Gamma^+(S_\ell)$. Thus, some event whose variables intersect with those in A must occur in the M run. But when the algorithm attempts to update variables shared by these two events in the M run, it will access values beyond the final round of the resampling table, a contradiction. □

Theorem 8. *When Condition 5 holds and Algorithm 2 halts, its output is a product distribution conditioned on avoiding all bad events.*

Proof. Clearly the output of Algorithm 2 avoids all bad events, since otherwise it does not halt.

Let an independent set sequence $\mathcal{S}$ of length ℓ be the log of any successful run. By Lemma 7, conditioned on log $\mathcal{S}$, the output assignment σ is a product distribution conditioned on $\bigwedge_{i \notin \Gamma^+(S_\ell)} \overline{A_i}$. Since the algorithm terminates we further condition on none of the other bad events occurring. The resulting output distribution still has the property that valid assignments occur with probability proportional to those in the product distribution. Since the above argument is valid for any possible log, the theorem follows. □

In other words, let S be the set of assignments that avoid all bad events. Let U be the output of Algorithm 2. In the case that all variables are sampled from the uniform distribution, we have $\Pr(U = s) = \frac{1}{|S|}$, for all $s \in S$.

3. Expected running time of Algorithm 2

We first give a combinatorial explanation of q_I for any independent set I of the dependency graph G. To simplify the notation, we denote the event $\bigwedge_{i\in S} A_i$, i.e., the conjunction of all events in S, by $A(S)$.

For any independent set I in the dependency graph, we denote by p_I the probability $\Pr(A(I)$ that all events in I happen (and possibly some other events too). Since I is an independent set, any two events in I are independent. Hence,

$$p_I = \prod_{i\in I} p_i. \tag{2}$$

Moreover, for any set J of events that is not an independent set, we have $p_J = 0$ due to Condition 5.

On the other hand, the quantity q_I is in fact the probability that exactly the events in I happen and no others do. This can be verified using inclusion-exclusion, together with Condition 5:

$$\begin{aligned}\Pr\left(\bigwedge_{i\in I} A_i \wedge \bigwedge_{i\notin I} \overline{A_i}\right) &= \sum_{J\supseteq I} (-1)^{|J\setminus I|} p_J \\ &= \sum_{J\in\mathcal{I},\ J\supseteq I} (-1)^{|J\setminus I|} p_J = q_I, \end{aligned} \tag{3}$$

where $\mathcal{I}$ denotes the collection of all independent sets of G. Since the events $(\bigwedge_{i\in I} A_i \wedge \bigwedge_{i\notin I} \overline{A_i})$ are mutually exclusive for different I's,

$$\sum_{I\in\mathcal{I}} q_I = 1.$$

Moreover, since the event $A(I)$ is the union over $J \supseteq I$ of the events $(\bigwedge_{i\in J} A_i \wedge \bigwedge_{i\notin J} \overline{A_i})$, we have

$$p_I = \sum_{J\in\mathcal{I},\ J\supseteq I} q_J. \tag{4}$$

Recall Definition 6 of independent set sequences. For an independent set sequence $\mathcal{S}$ of length $\ell > 0$, let

$$p_{\mathcal{S}} = \prod_{i=1}^{\ell} p_{S_i}. \tag{5}$$

Moreover, we follow the convention that $p_{\mathcal{S}} = 1$ if $\mathcal{S}$ is empty.

Let I be an independent set. We say an independent set sequence $\mathcal{S}$ of length ℓ is *compatible* with I if $I \subseteq \Gamma^+(S_\ell)$. When Algorithm 2 runs, the variable I goes through a sequence of values $I_1, I_2, \ldots$. We say that Algorithm 2 runs with log $\mathcal{S}$ up to round ℓ if $\mathcal{S} = (I_1, \ldots, I_\ell)$. The probability that Algorithm 2 runs with log $\mathcal{S}$ up to round ℓ and that $I_{\ell+1} = I$ has a simple expression.

Lemma 9. *Let $I \in \mathcal{I}$ and $\mathcal{S}$ be an independent set sequence of length ℓ that is compatible with I. Then*

$$\Pr\left(\textit{after } \ell \textit{ rounds the log is } \mathcal{S} \textit{ and } I_{\ell+1} = I\right) = q_I p_{\mathcal{S}}.$$

Proof. Clearly, if $q_I = 0$, then the said sequence will never happen. We assume that $q_I > 0$ in the following.

Let μ be the product distribution of sampling all variables uniformly and independently. We need to distinguish the probability space with respect to μ from that with respect to the execution of the algorithm. When we write $\Pr_{\mathrm{PRS}}(\cdot)$, it means with respect to the algorithm. When we write $\Pr_\mu(\cdot)$, it means the (static) space with respect to μ. As noted before, to simplify

the notation we will use $A(S)$ to denote the event $\bigwedge_{i\in S} A_i$, where $S \subseteq [m]$. In addition, $B(S)$ will be used to denote $\bigwedge_{i\in S} \overline{A_i}$. For $I \in \mathcal{I}$, define

$$\partial I = \Gamma^+(I) \setminus I, \quad I^e = [m] \setminus \Gamma^+(I), \quad \text{and} \quad I^c = [m] \setminus I = \partial I \cup I^e.$$

So ∂I is the "boundary" of I, comprising events that are not in I but which depend on I, and I^e is the "exterior" of I, comprising events that are independent of all events in I. The complement I^c is simply the set of all events not in I. Note that $B(I^c) = B(\partial I) \wedge B(I^e)$. As examples of the notation, $\Pr_\mu(A(I)) = \prod_{i\in I} p_i = p_I$ is the probability that all events in I occur under μ, and $\Pr_\mu(A(I) \wedge B([m] \setminus I)) = q_I$ is probability that exactly the events in I occur.

By the definition of I^e, we have that

$$\Pr\left(B(I^e) \mid A(I)\right) = \Pr(B(I^e)), \tag{6}$$

and, by Condition 5, that

$$A(I) \wedge B(\partial I) = A(I). \tag{7}$$

In the first round, since all variables are fresh, the distribution is just μ. Thus the probability of getting *exactly* S_1 in the first round is q_{S_1}, or, equivalently,

$$\begin{aligned} q_{S_1} &= \Pr_\mu\left(A(S_1) \wedge B(S_1^c)\right) \\ &= \Pr_\mu\left(A(S_1) \wedge B(\partial S_1) \wedge B(S_1^e)\right) \\ &= \Pr_\mu\left(A(S_1) \wedge B(S_1^e)\right) && \text{by (7)} \\ &= \Pr_\mu\left(A(S_1)\right) \Pr_\mu\left(B(S_1^e)\right). && \text{by (6)} \end{aligned}$$

Clearly the calculation above holds for an arbitrary independent set $I \in \mathcal{I}$. Hence,

$$q_I = \Pr_\mu\left(A(I)\right) \Pr_\mu\left(B(I^e)\right). \tag{8}$$

At round t, since we only resample variables that are involved in S_t, we have that $S_{t+1} \subseteq \Gamma^+(S_t)$. Moreover, variables are not resampled in any A_i where $i \notin \Gamma^+(S_t)$, and hence

$$B(S_{t+1}^c) \wedge B(S_t^e) = B(S_{t+1}^c). \tag{9}$$

Conditioned on S_t, by Lemma 7, the distribution of σ_{t+1} at round $t+1$ is uniform conditioned on none of the events outside of $\Gamma^+(S_t)$ occuring; namely, it is $\Pr_\mu(\cdot \mid B(S_t^e))$. Thus the probability of the event $I_{t+1} = S_{t+1}$ is

$$\begin{aligned} \Pr_{\mathrm{PRS}}\left(A(S_{t+1}) \wedge B(S_{t+1}^c)\right) &= \Pr_\mu\left(A(S_{t+1}) \wedge B(S_{t+1}^c) \mid B(S_t^e)\right) \\ &= \frac{\Pr_\mu\left(A(S_{t+1}) \wedge B(S_{t+1}^c) \wedge B(S_t^e)\right)}{\Pr_\mu(B(S_t^e))} \\ &= \frac{\Pr_\mu\left(A(S_{t+1}) \wedge B(S_{t+1}^c)\right)}{\Pr_\mu(B(S_t^e))} && \text{by (9)} \\ &= \frac{\Pr_\mu\left(A(S_{t+1})\right) \Pr_\mu\left(B(S_{t+1}^e)\right)}{\Pr_\mu(B(S_t^e))} && \text{by (8).} \end{aligned} \tag{10}$$

Recursively applying (10) we see that the probability of getting the log $\mathcal{S}$ and final state I is

$$\begin{aligned} \Pr_{\mathrm{PRS}}\left(\text{the log is } \mathcal{S} \text{ and } I_{\ell+1} = I\right) &= \Pr_\mu\left(A(S_1)\right) \Pr_\mu\left(B(S_1^e)\right) \prod_{t=2}^{\ell} \frac{\Pr_\mu\left(A(S_{t+1})\right) \Pr_\mu\left(B(S_{t+1}^e)\right)}{\Pr_\mu(B(S_t^e))} \\ &= \Pr_\mu\left(A(I)\right) \Pr_\mu\left(B(I^e)\right) \prod_{t=1}^{\ell} \Pr_\mu\left(A(S_t)\right) \\ &= q_I p_{\mathcal{S}}, \end{aligned}$$

where in the last line we used (8) and (5). □

Essentially the proof above is a delayed revelation argument. At each round $1 \le t \le \ell$, we only reveal variables that are involved in S_t. Thus, at round ℓ, we have revealed all variables that are involved in $\mathcal{S}$. With respect to these variables, the sequence $\mathcal{S}$ happens with probability $p_{\mathcal{S}}$. Condition 5 guarantees that what we have revealed so far does not interfere with the final output (cf. Lemma 7). Hence the final state happens with probability q_I.

In particular, Lemma 9 implies the following equality, which is first shown by Kolipaka and Szegedy [16] via a different method.

Corollary 10. *If $q_\emptyset > 0$, then*

$$\sum_{\mathcal{S} \text{ s.t. } S_1 = I} p_{\mathcal{S}} q_\emptyset = q_I,$$

where I is a non-empty independent set of G.

Proof. First we claim that if $q_\emptyset > 0$, then Algorithm 2 halts with probability 1. Indeed, if we resample a subset of variables, the probability of getting an assignment that avoids all events is at least $q_\emptyset$. Hence the probability that the algorithm does not halt at time t is at most $(1-q_\emptyset)^t$, which goes to 0 as t goes to infinity.

Then we apply Lemma 9 when $\emptyset$ is the final independent set. The left hand side is the total probability of all possible halting logs whose first independent set is exactly I. This is equivalent to getting exactly I in the first step, which happens with probability q_I. □

As a sanity check, the probability of all possible logs should sum to 1 when $q_\emptyset > 0$ and the algorithm halts with probability 1. Indeed, by Corollary 10,

$$\begin{aligned}\sum_{\mathcal{S}} p_{\mathcal{S}} q_\emptyset &= q_\emptyset + \sum_{I \in \mathcal{I} \text{ s.t. } I \neq \emptyset} \; \sum_{\mathcal{S} \text{ s.t. } S_1 = I} q_\emptyset p_{\mathcal{S}} \\ &= \sum_{I \in \mathcal{I}} q_I = 1,\end{aligned}$$

where $\mathcal{S}$ is an independent set sequence. In other words,

$$\sum_{\mathcal{S}} p_{\mathcal{S}} = \frac{1}{q_\emptyset}, \tag{11}$$

where $\mathcal{S}$ is an independent set sequence. This fact is also observed by Knuth [15, Page 86, Theorem F] and Harvey and Vondrák [13, Corollary 5.28]. Our proof here gives a combinatorial explanation of this equality.

Equation (11) holds whenever $q_\emptyset > 0$. Recall that $q_\emptyset$ is the shorthand of $q_\emptyset(\mathbf{p})$, which is

$$q_\emptyset(\mathbf{p}) = \sum_{I \in \mathcal{I}} (-1)^{|I|} \prod_{i \in I} p_i, \tag{12}$$

where $\mathcal{I}$ is the collection of independent sets of the dependency graph G.

Lemma 11. *If $q_\emptyset(\boldsymbol{p}) > 0$, then $q_\emptyset(p_1, \dots, p_i z, \dots, p_m) > 0$ for any $i \in [m]$ and $0 \le z \le 1$.*

Proof. By (12),

$$q_\emptyset(p_1, \dots, p_i z, \dots, p_m) = \sum_{I \in \mathcal{I},\, i \notin I} (-1)^{|I|} \prod_{j \in I} p_j + z \sum_{I \in \mathcal{I},\, i \in I} (-1)^{|I|} \prod_{j \in I} p_j.$$

Notice that $\sum_{I \in \mathcal{I},\, i \in I} (-1)^{|I|} \prod_{j \in I} p_j = -q_i(\mathbf{p}) < 0$ ($q_i(\mathbf{p})$ is the probability of exactly event A_i occurring). Hence $q_\emptyset(p_1, \dots, p_i z, \dots, p_m) \ge q_\emptyset(\mathbf{p}) > 0$. □

Let T_i be the number of resamplings of event A_i and T be the total number of resampling events. Then $T = \sum_{i=1}^m T_i$.

Lemma 12. *If $q_\emptyset(\boldsymbol{p}) > 0$, then $\mathbb{E}\, T_i = q_\emptyset(\boldsymbol{p}) \left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)} \right)' \Big|_{z=1}$.*

Proof. By Lemma 11, Equation (11) holds with p_i replaced by $p_i z$ where $z \in [0,1]$. For a given independent set sequence $\mathcal{S}$, let $T_i(\mathcal{S})$ be the total number of occurences of A_i in $\mathcal{S}$. Then we have that

$$\sum_{\mathcal{S}} p_{\mathcal{S}} z^{T_i(\mathcal{S})} = \frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}. \tag{13}$$

Take derivative with respect to z of (13):

$$\sum_{\mathcal{S}} T_i(\mathcal{S}) p_{\mathcal{S}} z^{T_i(\mathcal{S})-1} = \left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}\right)'.$$

Evaluate the equation above at $z = 1$:

$$\sum_{\mathcal{S}} T_i(\mathcal{S}) p_{\mathcal{S}} = \left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}\right)'\Big|_{z=1}. \tag{14}$$

On the other hand, we have that

$$\begin{aligned}
\mathbb{E}\, T_i &= \sum_{\mathcal{S}} \Pr_{\mathrm{PRS}}(\text{the log is } \mathcal{S})\, T_i(\mathcal{S}) \\
&= \sum_{\mathcal{S}} p_{\mathcal{S}} q_\emptyset(\mathbf{p}) T_i(\mathcal{S}) && \text{(by Lemma 9)} \\
&= q_\emptyset(\mathbf{p}) \left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}\right)'\Big|_{z=1}. && \text{(by (14))}
\end{aligned}$$

This completes the proof. □

Theorem 13. *If $q_\emptyset > 0$, then $\mathbb{E}\, T = \sum_{i=1}^m \frac{q_i}{q_\emptyset}$.*

Proof. Clearly $\mathbb{E}\, T = \sum_{i=1}^m \mathbb{E}\, T_i$. By Lemma 12, all we need to show is that

$$q_\emptyset(\mathbf{p}) \left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}\right)'\Big|_{z=1} = \frac{q_i(\mathbf{p})}{q_\emptyset(\mathbf{p})}. \tag{15}$$

This is because

$$\begin{aligned}
q'_\emptyset(p_1, \dots, p_i z, \dots, p_m) &= \sum_{i \in J,\, J \in \mathcal{I}} (-1)^{|J|} \prod_{j \in J} p_j \\
&= -q_i(\mathbf{p}),
\end{aligned}$$

and thus

$$\begin{aligned}
\left(\frac{1}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)}\right)' &= \frac{-q'_\emptyset(p_1, \dots, p_i z, \dots, p_m)}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)^2} \\
&= \frac{q_i(\mathbf{p})}{q_\emptyset(p_1, \dots, p_i z, \dots, p_m)^2}.
\end{aligned}$$

It is easy to see that (15) follows as we set $z = 1$ and the theorem is shown. □

The quantity $\sum_{i=1}^m \frac{q_i}{q_\emptyset}$ is not always easy to bound. It is well-known that when the probability vector $\mathbf{p}$ satisfies Shearer's condition with a constant "slack", the running time is in fact linear in the number of events. We can derive it using a result in [23], where Scott and Sokal characterized a region for the independence polynomial to be zero-free. We rewrite it in our terms.

Theorem 14 (Corollary 5.7 in [23])**.** *Let $d \geq 2$ be a positive integer and $p_c = \frac{(d-1)^{(d-1)}}{d^d}$. If G has maximum degree d, then $q_\emptyset(p_c \mathbf{1}) > 0$.*

Corollary 15. *Let $d \geq 2$ be a positive integer and $p_c = \frac{(d-1)^{(d-1)}}{d^d}$. Let $p = \max_{i \in [m]}\{p_i\}$. If G has maximum degree d and $p < p_c$, then $\mathbb{E}\, T \leq \frac{p}{p_c - p} \cdot m$.*

Proof. By Theorem 14,

$$q_\emptyset(p_c\mathbf{1}) > 0. \tag{16}$$

By (16) and Lemma 11, we have that for any $i \in [m]$,

$$q_\emptyset(p_1, \ldots, \delta p_i, \ldots, p_m) > 0,$$

where $\delta = \frac{p_c}{p} > 1$ so that $\delta p_i = p_c \cdot \frac{p_i}{p} \leq p_c$. By (12), we have that

$$\begin{aligned} 0 < q_\emptyset(p_1, \ldots, \delta p_i, \ldots, p_m) &= \sum_{I \in \mathcal{I},\, i \notin I} (-1)^{|I|} \prod_{j \in I} p_j + \delta \sum_{I \in \mathcal{I},\, i \in I} (-1)^{|I|} \prod_{j \in I} p_j \\ &= \sum_{I \in \mathcal{I},\, i \notin I} (-1)^{|I|} \prod_{j \in I} p_j - \delta q_i(\mathbf{p}). \end{aligned}$$

Or in other words,

$$\sum_{I \in \mathcal{I},\, i \notin I} (-1)^{|I|} \prod_{j \in I} p_j \geq \delta q_i(\mathbf{p}).$$

Thus by (12) again,

$$\begin{aligned} q_\emptyset(\mathbf{p}) &= \sum_{I \in \mathcal{I},\, i \notin I} (-1)^{|I|} \prod_{j \in I} p_j - q_i(\mathbf{p}) \\ &\geq \delta q_i(\mathbf{p}) - q_i(\mathbf{p}) = (\delta - 1) q_i(\mathbf{p}). \end{aligned}$$

Hence,

$$\frac{q_i(\mathbf{p})}{q_\emptyset(\mathbf{p})} \leq \frac{1}{\delta - 1}.$$

Due to Theorem 13,

$$\mathbb{E}\, T = \sum_{i=1}^{m} \frac{q_i(\mathbf{p})}{q_\emptyset(\mathbf{p})} \leq \frac{m}{\delta - 1}. \qquad \square$$

4. Applications of Algorithm 2

4.1. **Sink-free Orientations.** The goal of this application is to sample a sink-free orientation. Given a graph $G = (V, E)$, an orientation of edges is a mapping σ so that $\sigma(e) = (u, v)$ or (v, u) where $e = (u, v) \in E$. A *sink* under orientation σ is a vertex v so that for any adjacent edge $e = (u, v)$, $\sigma(e) = (u, v)$. A sink-free orientation is an orientation so that no vertex is a sink.

Name: Sampling Sink-free Orientations
Instance: A Graph G.
Output: A uniform sink-free orientation.

The first algorithm for this problem is given by Bubley and Dyer [6], using Markov chains and path coupling techniques.

In this application, we associate with each edge a random variable, whose possible values are (u, v) or (v, u). For each vertex v_i, we associate it with a bad event A_i, which happens when v_i is a sink. Thus the graph G itself is also the dependency graph. Condition 5 is satisfied. This follows because if a vertex is a sink, then none of its neighbours can be a sink. Thus we may apply Algorithm 2, which yields Algorithm 3. This is the "sink-popping" algorithm given by Cohn, Pemantle, and Propp [7].

Algorithm 3 Sample Sink-free Orientations

(1) Orient each edge independently and uniformly at random.

(2) While there is at least one sink, re-orient all edges that are adjacent to a sink.

(3) Output the current assignment.

Let $Z_{\text{sink},0}$ be the number of sink-free orientations, and let $Z_{\text{sink},1}$ be the number of orientations having exactly one sink. Then Theorem 13 specializes into the following.

Theorem 16. *The expected number of resampled sinks in Algorithm 3 is* $\frac{Z_{\text{sink},1}}{Z_{\text{sink},0}}$.

The next theorem gives an explicit bound on $\frac{Z_{\text{sink},1}}{Z_{\text{sink},0}}$.

Theorem 17. *Let G be a connected graph on n vertices. If G is not a tree, then* $\frac{Z_{\text{sink},1}}{Z_{\text{sink},0}} \leq n(n-1)$, *where $n = |V(G)|$.*

Proof. Consider an orientation of the edges of G with a unique sink at vertex v. We give a systematic procedure for transforming this orientation to a sink-free orientation. Since G is connected and not a tree, there exists an (undirected) path Π in G of the form $v = v_0, v_1, \ldots, v_{\ell-1}, v_\ell = v_k$, where the vertices $v_0, v_1, \ldots, v_{\ell-1}$ are all distinct and $0 \leq k \leq \ell - 2$. In other words, Π is a simple path of length $\ell - 1$ followed by a single edge back to some previously visited vertex. We will choose a canonical path of this form (depending only on G and not on the current orientation) for each start vertex v.

We now proceed as follows. Since v is a sink, the first edge on Π is directed (v_1, v_0). Reverse the orientation of this edge so that it is now oriented (v_0, v_1). This operation destroys the sink at $v = v_0$ but may create a new sink at v_1. If v_1 is not a sink then halt. Otherwise, reverse the orientation of the second edge of Π from (v_2, v_1) to (v_1, v_2). Continue in this fashion: if we reach v_i and it is not a sink then halt; otherwise reverse the orientation of the $(i+1)$th edge from (v_{i+1}, v_i) to (v_i, v_{i+1}). This procedure must terminate with a sink-free graph before we reach v_ℓ. To see this, note that if we reach the vertex $v_{\ell-1}$ then the final edge of Π must be oriented $(v_{\ell-1}, v_\ell)$, otherwise the procedure would have terminated already at vertex $v_k (= v_\ell)$.

The effect of the above procedure is to reverse the orientation of edges on some initial segment $v_0, \ldots, v_i$ of Π. To put the procedure into reverse, we just need to know the identity of the vertex v_i. So our procedure associates at most n orientations having a single sink at vertex v with each sink-free orientation. There are $n(n-1)$ choices for the pair (v, v_i), and hence $n(n-1)$ single-sink orientations associated with each sink-free orientation. This establishes the result. □

Remark. *The bound in Theorem 17 is optimal up to a factor of 2. Consider a cycle of length n. Then there are 2 sink-free orientations, and $n(n-1)$ single-sink orientations.*

Theorem 17 and Theorem 16 together yield a n^2 bound on the expected number of resamplings that occur during a run of Algorithm 3. A cycle of length n is an interesting special case. Consider the number of clockwise oriented edges during a run of the algorithm. It is easy to check that this number evolves as an unbiased lazy simple random walk on $[0, n]$. Since the walk starts close to $n/2$ with high probability, we know that it will take $\Omega(n^2)$ steps to reach one of the sink-free states, i.e., 0 or n.

On the other hand, if G is a regular graph of degree $\Delta \geq 3$, then we get a much better linear bound from Corollary 15. In the case $\Delta = 3$, we have $p_c = 4/27$, $p = 1/8$ and $p/(p_c - p) = 27/5$. So the expected number of resamplings is bounded by $27n/5$. The constant in the bound improves as Δ increases. Conversely, since the expected running time is exact, we can also apply Corollary 15 to give an upper bound of $\frac{Z_{sink,1}}{Z_{sink,0}}$ when G is a regular graph.

4.2. **Rooted Spanning Trees.** Given a graph $G = (V, E)$ with a special vertex r, we want to sample a uniform spanning tree with r as the root.

Name: Sampling Rooted Spanning Trees
Instance: A Graph G with a vertex r.
Output: A uniform spanning tree rooted at r.

Of course, any given spanning tree may be rooted at any vertex r, so there is no real difference between rooted and unrooted spanning trees. However, since this approach to sampling generates an oriented tree, it is easier to think of the trees as being rooted at a particular vertex r.

For all vertices other than r, we randomly assign it to point to one of its neighbours. This is the random variable associated with v. We will think of this random variable as an arrow $v \to s(v)$ pointing from v to its successor $s(v)$. The arrows point out an oriented subgraph of G with $n-1$ edges $\{\{v, s(v)\} : v \in V \setminus \{r\}\}$ directed as specified by the arrows. The constraint for this subgraph to be a tree rooted at r is that it contains no directed cycles. Note that there are $2^{|E|-|V|+\kappa(G)}$ (undirected) cycles in G, where $\kappa(G)$ is the number of connected components of G. Hence, we have possibly exponentially many constraints.

Two cycles are dependent if they share at least one vertex. We claim that Condition 5 is satisfied. Suppose a cycle C is present, and $C' \neq C$ is another cycle that shares at least one vertex with C. If C' is also present, then we may start from any vertex $v \in C \cap C'$, and then follow the arrows $v \to v'$. Since both C and C' are present, it must be that $v' \in C \cap C'$ as well. Continuing this argument we see that $C = C'$. Contradiction!

As Condition 5 is met, we may apply Algorithm 2, yielding Algorithm 4. This is exactly the "cycle-popping" algorithm by Wilson [28], as described in [22].

Algorithm 4 Sample Rooted Spanning Trees

(1) Let T be an empty set. For each vertex $v \neq r$, randomly choose a neighbour $u \in \Gamma(v)$ and add an edge (v, u) to T.

(2) While there is at least one cycle in T, remove all edges in all cycles, and for all vertices whose edges are removed, redo step (1).

(3) Output the current set of edges.

Let $Z_{\text{tree},0}$ be the number of possible assignments of arrows to the vertices of G, that yield a (directed) tree with root r, and $Z_{\text{tree},1}$ be the number of assignments that yield a unicyclic subgraph. A unicyclic subgraph has two components: a directed tree with root r, and a directed cycle with a number of directed subtrees rooted on the cycle. The next theorem gives an explicit bound on $\frac{Z_{\text{tree},1}}{Z_{\text{tree},0}}$.

Theorem 18. *Suppose G is a connected graph on n vertices, with m edges. Then $\frac{Z_{\text{tree},1}}{Z_{\text{tree},0}} \leq mn$.*

Proof. Consider a assignment of arrows to the vertices of G that forms a unicyclic graph. As previously observed, this graph has two components. As G is connected, there must be an edge in G joining the two components; let this edge be $\{v_0, v_1\}$, where v_0 is in the tree component and v_1 in the unicyclic component. Now extend this edge to a path $v_0, v_1, \ldots, v_\ell$, by following arrows until we reach the cycle. Thus, $v_1 \to v_2$, $v_2 \to v_3$, $\ldots$, $v_{\ell-1} \to v_\ell$ are all arrows, and v_ℓ is the first vertex that lies on the cycle. (It may happen that $\ell = 1$.) Let $v_\ell \to v_{\ell+1}$ be the arrow out of v_ℓ. Now reassign the arrows from vertices $v_1, \ldots, v_\ell$ thus: $v_\ell \to v_{\ell-1}$, $\ldots$, $v_2 \to v_1$, $v_1 \to v_0$. Notice that the result is a directed tree rooted at r.

As before, we would like to bound the number of unicyclic subgraphs associated with a given tree by this procedure. We claim that the procedure can be reversed given just two pieces of information, namely, the edge $\{v_\ell, v_{\ell+1}\}$ and the vertex v_0. Note that, even though the edge $\{v_\ell, v_{\ell+1}\}$ is undirected, we can disambiguate the endpoints, as v_ℓ is the vertex closer to the root r. The vertices $v_{\ell-1}, \ldots, v_1$ are easy to recover, as they are the vertices on the unique path in the tree from v_ℓ to v_0. To recover the unicyclic subgraph, we just need to reassign the arrows for vertices $v_1, \ldots, v_\ell$ as follows: $v_1 \to v_2$, $\ldots$, $v_\ell \to v_{\ell+1}$.

As the procedure can be reversed knowing one edge and one vertex, the number of unicyclic graphs associated with each tree can be at most mn. □

Theorem 18 combined with Theorem 13 yields an mn upper bound on the expected number of "popped cycles" during a run of Algorithm 4.

On the other hand, the bound of Theorem 18 is tight up to constant factors. For example, take a cycle of length n. Then there are n spanning trees with a particular root v, and there are $\Omega(n^3)$ unicyclic graphs (here a cycle has to be of length 2). Thus the ratio is $\Omega(n^2) = \Omega(mn)$ since $m = n$.

4.3. **Extremal CNF formulas.** A classical setting in the study of algorithmic Lovász Local Lemma is to find satisfying assignments in k-CNF formulas[2], when the number of appearances of every variable is bounded by d. Theorem 1 guarantees the existence of a satisfying assignment as long as $d \leq \frac{2^k}{ek} + 1$. On the other hand, sampling is apparently harder than searching in this setting. As shown in [4, Corollary 30], it is **NP**-hard to approximately sample satisfying assignments when $d \geq 5 \cdot 2^{k/2}$, even restricted to the special case of monotone formulas.

Meanwhile, sink-free orientations can be recast in terms of CNF formulas. Every vertex in the graph is mapped to a clause, and every edge is a variable. Thus every variable appears exactly twice, and we require that the two literals of the same variable are always opposite. Interpret an orientation from u to v as making the literal in the clause corresponding to v false. The "sink-free" requirement is thus "not all literals in a clause are false". Hence a "sink-free" orientation is just a satisfying assignment for the corresponding CNF formula.

To apply Algorithm 2, we need to require that the CNF formula satisfies Condition 5. Such formulas are defined as follows.

Definition 19. *We call a CNF formula* extremal *if for every two clauses C_i and C_j, if there is a common variable shared by C_i and C_j, then there exists some variable x such that x appears in both C_i and C_j and the two literals are one positive and one negative.*

Let $C_1, \dots, C_m$ be the clauses of a formula φ. Then define the bad event A_i as the unsatisfying assignment of clause C_i. For an extremal CNF formula, these bad events satisfy Condition 5. This is because if $A_i \sim A_j$, then by Definition 19, there exists a variable $x \in \mathsf{var}(A_i) \cap \mathsf{var}(A_j)$ such that the unsatisfying assignment of C_i and C_j differ on x. Hence $A_i \cap A_j = \emptyset$.

In this formulation, if the size of C_i is k, then the corresponding event A_i happens with probability $p_i = \Pr(A_i) = 2^{-k}$. Suppose each variable appears at most d times. Then the maximum degree in the dependency graph is at most $\Delta = (d-1)k$. Note that in Corollary 15, $p_c = \frac{(\Delta-1)^{(\Delta-1)}}{\Delta^\Delta} \geq \frac{1}{e\Delta}$. Thus if $d \leq \frac{2^k}{ek} + 1$, then $p_i = 2^{-k} < p_c$ and we may apply Corollary 15 to obtain a polynomial time sampling algorithm.

Corollary 20. *For extremal k-CNF formulas where each variable appears in at most d clauses, if $d \leq \frac{2^k}{ek} + 1$, then Algorithm 2 samples satisfying assignments uniformly at random, with $O(m)$ expected resamplings where m is the number of clauses.*

The condition in Corollary 20 essentially matches the condition of Theorem 1. On the other hand, let $Z_{CNF,0}$ be the number of satisfying assignments, and $Z_{CNF,1}$ be the number of assignments satisfying all but one clause. If we only require Shearer's condition in Theorem 2, then the expected number of resamplings $\frac{Z_{CNF,1}}{Z_{CNF,0}}$ can be exponential, as shown in the next example.

Example. *Construct an extremal CNF formula $\varphi = C_1 \wedge C_2 \wedge \cdots \wedge C_{4m}$ as follows. Let $C_1 := x_1$. Then the variable x_1 is pinned to 1 to satisfy C_1. Let $C_2 := \overline{x}_1 \vee y_1 \vee y_2$, $C_3 := \overline{x}_1 \vee y_1 \vee \overline{y}_2$, and $C_4 := \overline{x}_1 \vee \overline{y}_1 \vee y_2$. Then y_1 and y_2 are also pinned to 1 to satisfy all $C_1 - C_4$.*

We continue this construction by letting

$$
\begin{aligned}
C_{4k+1} &:= \overline{y}_{2k-1} \vee \overline{y}_{2k} \vee x_{k+1},\\
C_{4k+2} &:= \overline{x}_{k+1} \vee y_{2k+1} \vee y_{2k+2},\\
C_{4k+3} &:= \overline{x}_{k+1} \vee y_{2k+1} \vee \overline{y}_{2k+2},\\
C_{4k+4} &:= \overline{x}_{k+1} \vee \overline{y}_{2k+1} \vee y_{2k+2},
\end{aligned}
$$

for all $1 \leq k \leq m-1$. It is easy to see by induction that to satisfy all of them, all x_i's and y_i's have to be 1. Moreover, one can verify that this is indeed an extremal formula. Thus $Z_{CNF,0} = 1$.

On the other hand, if we are allowed to ignore C_1, then x_1 can be 0. In that case, there are 3 choices of y_1 and y_2 so that x_2 to be 0 as well. Thus, there are at least 3^m assignments that only violate C_1, where $x_1 = x_2 = \cdots = x_m = 0$. It implies that $Z_{CNF,1} \geq 3^m$. Hence we see that

[2]As usual in the study of Lovász Local Lemma, by "k-CNF" we mean that every clause has exactly size k.

$\frac{Z_{CNF,1}}{Z_{CNF,0}} \geq 3^m$. *Due to Theorem 13, the expected running time of Algorithm 2 on this formula φ is at least exponential in m.*

We will discuss more about sampling satisfying assignments of a k-CNF formula in Section 7.1.

5. Generalized Partial Rejection Sampling

In this section we give a more general version of Algorithm 2 which can be applied to arbitrary instances in the variable framework, even without Condition 5.

Let $X = \{x_1, \ldots, x_n\}$ be the set of random variables. Each X_i can have its own distribution and range D_i. Let $A_1, \ldots, A_m$ be the "bad" events and $G = (V, E)$ be the dependency graph. Let σ be an assignment of all variables so that $\sigma(x_i) \in D_i$. (Note that we allow the range of x_i to be dependent on i.) Given σ, let $\mathsf{Bad}(\sigma)$ be the set of occurring bad events; that is, $\mathsf{Bad}(\sigma) = \{i \mid \sigma \in A_i\}$. For a subset $S \subset V$, let ∂S be the boundary of S; that is, $\partial S = \{i \mid i \notin S \text{ and } \exists j \in S,\ (i,j) \in E\}$. Moreover, let

$$\mathsf{var}(S) := \bigcup_{i \in S} \mathsf{var}(A_i).$$

Let $\sigma|_S$ (or simply σ_S when there is no confusion) be the partial assignment of σ restricted to $\mathsf{var}(S)$. For an event A_i and $S \subseteq V$, we write $A_i \cap \sigma_S = \emptyset$ if $\mathsf{var}(A_i) \cap \mathsf{var}(S) = \emptyset$, or A_i is disjoint from the partial assignment of σ, when both are restricted to $\mathsf{var}(A_i) \cap \mathsf{var}(S)$. Informally, A_i cannot occur given partial assignment σ_S. Otherwise we write $A_i \cap \sigma_S \neq \emptyset$.

Definition 21. *A set $S \subseteq V$ is* unblocking *under σ if for every $i \in \partial S$, $A_i \cap \sigma_S = \emptyset$.*

Given σ, our goal is to resample a set of events that is unblocking and contains $\mathsf{Bad}(\sigma)$. Such a set must exist because V is unblocking (∂V is empty) and $\mathsf{Bad}(\sigma) \subseteq V$. However, we want to resample as few events as possible.

Algorithm 5 Select a resampling set $\mathsf{Res}(\sigma)$ under an assignment σ

(1) Let $R = \mathsf{Bad}(\sigma)$, which is the set of events that will be resampled. Let $N = \emptyset$, which is the set of events that will not be resampled.

(2) While $\partial R \setminus N \neq \emptyset$, go through $i \in \partial R \setminus N$; if $A_i \cap \sigma_R \neq \emptyset$, add i into R, otherwise add i into N.

(3) Output R.

Intuitively, we start by setting the resampling set R_0 as the set of bad events $\mathsf{Bad}(\sigma)$. We mark resampling events in rounds, similar to a breadth first search. Let R_t be the resampling set of round $t \geq 0$. In round $t+1$, let A_i be an event on the boundary of R_t that hasn't been marked yet. We mark it "resampling" if the partial assignments on the shared variables of A_i and R_t can be extended so that A_i occurs. Otherwise we mark it "not resampling". We continue this process until there is no unmarked event left on the boundary of the current R. An event outside of $\Gamma^+(R)$ may be left unmarked at the end of Algorithm 5. Note that once we mark some event "not resampling", it will never be added into the resampling set.

It it easy to see that the output of Algorithm 5 is deterministic under σ. Call it $\mathsf{Res}(\sigma)$.

Lemma 22. *Let σ be an assignment. For any $i \in \partial\mathsf{Res}(\sigma)$, $A_i \cap \sigma_{\mathsf{Res}(\sigma)} = \emptyset$.*

Proof. Since $i \in \partial\mathsf{Res}(\sigma)$, it must have been marked. Moreover, $i \notin \mathsf{Res}(\sigma)$, so it must be marked as "not resampling". Thus, there exists an intermediate set $R \subseteq \mathsf{Res}(\sigma)$ during the execution of Algorithm 5 such that $A_i \cap \sigma_R = \emptyset$ and $i \in \partial R$. It implies that A_i is disjoint from the partial assignment of σ restricted to $\mathsf{var}(A_i) \cap \mathsf{var}(R)$. However,

$$\mathsf{var}(A_i) \cap \mathsf{var}(R) \subseteq \mathsf{var}(A_i) \cap \mathsf{var}(\mathsf{Res}(\sigma))$$

as $R \subseteq \mathsf{Res}(\sigma)$. We have that $A_i \cap \sigma_{\mathsf{Res}(\sigma)} = \emptyset$. □

If Condition 5 is met, then $\mathsf{Res}(\sigma) = \mathsf{Bad}(\sigma)$. This is because at the first step, $R = \mathsf{Bad}(\sigma)$. By Condition 5, for any $i \in \partial\mathsf{Bad}(\sigma)$, A_i is disjoint from all A_j's where $j \in \mathsf{Bad}(\sigma)$ and $A_i \sim A_j$. Since A_j occurs under σ, $A_i \cap \sigma_R = \emptyset$. Algorithm 5 halts at the first step. In this case, since the resampling set is just the (independent) set of occurring bad events, the later Algorithm 6 coincides with Algorithm 2.

The key to $\mathsf{Res}(\sigma)$ is that if we change the values of variables outside of $\mathsf{Res}(\sigma)$, then $\mathsf{Res}(\sigma)$ does not change, unless a new bad event occur outside of $\mathsf{Res}(\sigma)$.

Lemma 23. *Let σ be an assignment. Let σ' be another assignment so that $\mathsf{Bad}(\sigma') \subseteq \mathsf{Res}(\sigma)$ and if $\sigma(x) \neq \sigma'(x)$, then $x \notin \mathsf{var}(A_i)$ for all $i \in \mathsf{Res}(\sigma)$. Then, $\mathsf{Res}(\sigma') = \mathsf{Res}(\sigma)$.*

Proof. Let $R_t(\sigma), N_t(\sigma)$ be the intermediate sets R, N, respectively, at time t of the execution of Algorithm 5 under σ. Thus $R_0(\sigma) = \mathsf{Bad}(\sigma)$ and $R_0(\sigma) \subseteq R_1(\sigma) \subseteq \cdots \subseteq \mathsf{Res}(\sigma)$. Moreover, $N_0(\sigma) \subseteq N_1(\sigma) \subseteq \cdots$. We will show by induction that $R_t(\sigma) = R_t(\sigma')$ and $N_t(\sigma) = N_t(\sigma')$ for any $t \geq 0$.

For the base case of $t = 0$, by the condition of the lemma, for every $i \in \mathsf{Bad}(\sigma) \subseteq \mathsf{Res}(\sigma)$, the assignments σ and σ' agree on $\mathsf{var}(A_i)$; or equivalently $\sigma_{\mathsf{Res}(\sigma)} = \sigma'_{\mathsf{Res}(\sigma)}$. Together with $\mathsf{Bad}(\sigma') \subseteq \mathsf{Res}(\sigma)$, it implies that $\mathsf{Bad}(\sigma) = \mathsf{Bad}(\sigma')$ and $R_0(\sigma) = R_0(\sigma')$. Moreover, $N_0(\sigma) = N_0(\sigma') = \emptyset$.

For the induction step $t > 0$, we have that $R_{t-1}(\sigma) = R_{t-1}(\sigma') \subseteq \mathsf{Res}(\sigma)$ and $N_{t-1}(\sigma) = N_{t-1}(\sigma')$. Let $R = R_{t-1}(\sigma) = R_{t-1}(\sigma')$ and $N = N_{t-1}(\sigma) = N_{t-1}(\sigma')$. Then we will go through $\partial R \setminus N$, which is the same for both σ and σ'. Moreover, while marking individual events "resampling" or not, it is sufficient to look at only $\sigma_R = \sigma'_R$ since $R \subseteq \mathsf{Res}(\sigma)$. Thus the markings are exactly the same, implying that $R_t(\sigma) = R_t(\sigma') \subseteq \mathsf{Res}(\sigma)$ and $N_t(\sigma) = N_t(\sigma')$. □

Algorithm 6 Generalized Partial Rejection Sampling

(1) Draw independent samples of all variables $X_1, \ldots, X_n$ from their respective distributions.

(2) While at least one bad event occurs under the current assignment σ, use Algorithm 5 to find $\mathsf{Res}(\sigma)$. Resample all variables in $\bigcup_{i \in \mathsf{Res}(\sigma)} \mathsf{var}(A_i)$.

(3) When none of the bad events holds, output the current assignment.

To prove the correctness of Algorithm 6, we will only use three properties of $\mathsf{Res}(\sigma)$, which are intuitively summarized as follows:

(1) $\mathsf{Bad}(\sigma) \subseteq \mathsf{Res}(\sigma)$;
(2) For any $i \in \partial\mathsf{Res}(\sigma)$, A_i is disjoint from the partial assignment of σ projected on $\mathsf{var}(A_i) \cap \mathsf{var}(\mathsf{Res}(\sigma))$ (Lemma 22);
(3) If we fix the partial assignment of σ projected on $\mathsf{var}(\mathsf{Res}(\sigma))$, then the output of Algorithm 5 is fixed, unless there are new bad events occurring outside of $\mathsf{Res}(\sigma)$ (Lemma 23).

Similarly to the analysis of Algorithm 2, we call $\mathcal{S} = S_1, \ldots, S_\ell$ the log, if S_i is the set of resampling events in step i of Algorithm 6. Note that for Algorithm 6, the log is not necessarily an independent set sequence. Also, recall that σ_i is the assignment of variables in step i.

Lemma 24. *Given any log $\mathcal{S}$ of length $\ell \geq 1$, $\sigma_{\ell+1}$ has the product distribution conditioned on none of A_i's occurring where $i \notin \Gamma^+(S_\ell)$.*

Proof. Suppose $i \notin \Gamma^+(S_\ell)$. By construction, S_ℓ contains all occurring bad events of σ_ℓ, and hence A_i does not occur under σ_ℓ. In step ℓ, we only resample variables that are involved in S_ℓ, so $\sigma_{\ell+1}$ and σ_ℓ agree on $\mathsf{var}(A_i)$. Hence A_i cannot occur under $\sigma_{\ell+1}$. Call an assignment σ *valid* if none of A_i occurs where $i \notin \Gamma^+(S_\ell)$. To show that $\sigma_{\ell+1}$ has the desired conditional product distribution, we will show that the probabilities of getting any two valid assignments σ and σ' are proportional to their probabilities of occurrence in the product distribution.

Let M be the resampling table so that the log of Algorithm 6 is $\mathcal{S}$ up to round ℓ, and $\sigma_{\ell+1} = \sigma$. Indeed, since we only care about events up to round $\ell + 1$, we may truncate the table so that

$M = \{X_{i,j} \mid 1 \le i \le n,\ \ 1 \le j \le i_{\ell+1}\}$. Let $M' = \{X'_{i,j} \mid 1 \le i \le n,\ \ 1 \le j \le i_{\ell+1}\}$ be another table where $X'_{i,j} = X_{i,j}$ if $j < i_{\ell+1}$ for any $i \in [n]$, but $X'_{i,i_{\ell+1}} = \sigma'(X_i)$. In other words, we only change the last assignment $(X_{i,i_{\ell+1}})$ to another valid assignment.

The lemma follows if Algorithm 6 running on M' generates the same log $\mathcal{S}$ up to round ℓ, since, if this is the case, then conditioned on the log $\mathcal{S}$, every possible table M where $\sigma_{\ell+1} = \sigma$ is one-to-one correspondence with another table M' where $\sigma'_{\ell+1} = \sigma'$. Hence the probability of getting σ is proportional to its weight in the product distribution.

Suppose otherwise and the log of running Algorithm 6 on M and M' differ. Let $t_0 \le \ell$ be the first round where resampling changes, by which we mean that $\mathsf{Res}(\sigma_{t_0}) \neq \mathsf{Res}(\sigma'_{t_0})$. By Lemma 23, there must be a variable X_i such that $i_{t_0} = i_{\ell+1}$ (otherwise $X_{i,i_{t_0}} = X'_{i,i_{t_0}}$) and X_i is involved in some event of $\mathsf{Res}(\sigma_{t_0})$ or X_i is involved in $\mathsf{Bad}(\sigma'_{t_0}) \setminus \mathsf{Res}(\sigma_{t_0})$.

(1) If X_i is involved in some event in $\mathsf{Res}(\sigma_{t_0})$, then X_i is resampled once more in the original run on M, and its index goes up to at least $i_{\ell+1}+1$ at round $\ell+1$. Contradiction!

This in particular implies that restricted to variables of $\mathsf{Res}(\sigma_{t_0})$, σ and σ' should agree; that is, $\sigma_{t_0}|_{\mathsf{Res}(\sigma_{t_0})} = \sigma'_{t_0}|_{\mathsf{Res}(\sigma_{t_0})}$.

(2) Otherwise there exists j such that $X_i \in \mathsf{var}(A_j)$, $j \in \mathsf{Bad}(\sigma'_{t_0})$ but $j \notin \mathsf{Res}(\sigma_{t_0})$.

Suppose first that $\forall k \in \mathsf{var}(A_j)$, $k_{t_0} = k_{\ell+1}$, which means that all variables of A_j have reached their final values in the M run at time t_0. This implies that $j \notin \Gamma^+(S_t)$ for any $t \ge t_0$ as otherwise some of the variables in $\mathsf{var}(A_j)$ would be resampled at least once after round t_0. In particular, $j \notin \Gamma^+(S_\ell)$. This contradicts with σ' being valid.

Otherwise there are some variables in $\mathsf{var}(A_j)$ that get resampled after time t_0 in the M run. Let t_1 be the first such time and $Y \subset \mathsf{var}(A_j)$ be the set of variables resampled at round t_1; namely, $Y = \mathsf{var}(A_j) \cap \mathsf{var}(\mathsf{Res}(\sigma_{t_1}))$. We have that $\sigma_{t_1}|_Y = \sigma_{t_0}|_Y$ because t_1 is the first time of resampling variables in Y. Moreover, as variables of Y have not reached their final values yet in the M run, $\sigma_{t_0}|_Y = \sigma'_{t_0}|_Y$. Thus, $\sigma_{t_1}|_Y = \sigma'_{t_0}|_Y$.

Assuming $j \in \mathsf{Res}(\sigma_{t_1})$ would contradict the fact that X_i has reached its final value in the M run. Hence $j \notin \mathsf{Res}(\sigma_{t_1})$, but nevertheless variables in $Y \subset \mathsf{var}(A_j)$ are resampled. This implies that $j \in \partial\mathsf{Res}(\sigma_{t_1})$. By Lemma 22, $A_j \cap \sigma_{t_1}|_{\mathsf{Res}(\sigma_{t_1})} = \emptyset$. This means that A_j is disjoint from the partial assignment of σ_{t_1} restricted to $\mathsf{var}(A_j) \cap \mathsf{var}(\mathsf{Res}(\sigma_{t_1})) = Y$. Equivalently, $A_j \cap \sigma_{t_1}|_Y = \emptyset$. However we know that $\sigma_{t_1}|_Y = \sigma'_{t_0}|_Y$, so $A_j \cap \sigma'_{t_0}|_Y = \emptyset$, contradicting $j \in \mathsf{Bad}(\sigma'_{t_0})$. □

Theorem 25. *If Algorithm 6 halts, then its output has the product distribution conditioned on none of A_i's occurring.*

Proof. Clearly the output of Algorithm 6 avoids all bad events, since otherwise it does not halt.

Let a sequence $\mathcal{S}$ of sets of events be the log of any successful run. By Lemma 24, conditioned on $\mathcal{S}$, the output assignment σ has the product distribution conditioned on $\bigwedge_{i \notin \Gamma^+(S_\ell)} \overline{A_i}$. Since the algorithm terminates we further condition on none of the other bad events occurring either. The resulting output distribution still has the property that valid assignments occur with probability proportional to those in the product distribution. Since the above argument is valid for any possible log, the theorem follows. □

6. Running Time Analysis of Algorithm 6

Obviously when there is no assignment avoiding all bad events, then Algorithm 6 will never halt. Thus we want to assume some conditions to guarantee a desired assignment. However, the optimal condition of Theorem 2 is quite difficult to work under. Instead, in this section we will be working under the assumption that Theorem 1's condition (1) holds. In fact, to make the presentation clean, we will mostly work with the simpler symmetric case.

However, as mentioned in Section 4.3, [4, Corollary 30] showed that even under the condition (1), sampling can still be **NP**-hard. We thus in turn look for some further condition to make Algorithm 6 efficient.

Let $\mu(\cdot)$ be the product distribution of sampling all variables independently. For two distinct events $A_i \sim A_j$, let R_{ij} be the event that the partial assignments on $\mathsf{var}(A_i) \cap \mathsf{var}(A_j)$ can be

extended to an assignment making A_j true. Thus, if A_i is added by Algorithm 5 at round $t \geq 1$, then for any event A_j added in round $t-1$ such that $A_i \sim A_j$, the event R_{ji} has to be true. Conversely, suppose A_i is unmarked and is under examination at round t. Then even if R_{ji} is true for all $j \in \Gamma(i)$ where A_j is added in round $t-1$, A_i is not necessarily marked "resampling" in round t by Algorithm 5. (It is possible for $A_i \cap \sigma_R = \emptyset$ even if all the R_{ji} are true.) Note that R_{ij} is not necessarily the same as R_{ji}. Let $r_{ij} := \mu(R_{ij})$.

Define $p := \max_{i \in [m]} p_i$ and $r := \max_{A_i \sim A_j,\, i \neq j} r_{ij}$. Let Δ be the maximum degree of the dependency graph G. The main result of the section is the following theorem.

Theorem 26. *Let m be the number of events and n be the number of variables. For any $\Delta \geq 2$, if $6ep\Delta^2 \leq 1$ and $3er\Delta < 1$, then the expected number of resampled events of Algorithm 6 is $O(m)$.*

Moreover, in this case, the expected number of rounds is $O(\log m)$ and thus the expected number of variable resamples is at most $O(n \log m)$.

The first condition $6ep\Delta^2 \leq 1$ is stronger than the condition of the symmetric Lovász Local Lemma, but this seems necessary since [4, Corollary 30] implies that if $p\Delta^2 \geq C$ for some constant C then the sampling problem is **NP**-hard. Intuitively, the second condition $3er\Delta < 1$ bounds the expansion from bad events to resampling events at every step of Algorithm 6. We will prove Theorem 26 in the rest of the section.

Let $S \subseteq [m]$ be a subset of vertices of the dependency graph G. Recall that $A(S)$ is the event $\bigwedge_{i \in S} A_i$ and $B(S)$ is the event $\bigwedge_{i \in S} \overline{A_i}$. Moreover, S^c is the complement of S, namely $S^c = [m] \setminus S$, and S^e is the "exterior" of S, namely $S^e = [m] \setminus \Gamma^+(S)$.

Lemma 24 implies that if we resample S at some step t of Algorithm 6, then at step $t+1$ the distribution is the product measure μ conditioned on none of the events in the exterior of S holds; namely $\Pr_\mu(\cdot \mid B(S^e))$.

Let E be an event (not necessarily one of A_i) depending on variables in $\mathsf{var}(E)$. Let $\Gamma(E) := \{i \mid i \in [m],\ \mathsf{var}(A_i) \cap \mathsf{var}(E) \neq \emptyset\}$ if E is not one of A_i, and $\Gamma(A_i) := \{j \mid j \in [m],\ j \neq i \text{ and } \mathsf{var}(A_j) \cap \mathsf{var}(A_i) \neq \emptyset\}$ is defined as usual. Let $S \subseteq [m]$ be a subset of vertices of G. The next lemma bounds the probability of E conditioned on none of the events in S happening.

Lemma 27. *Suppose (1) holds. For an event E and any set $S \subseteq [m]$,*

$$\Pr_\mu(E \mid B(S)) \leq \Pr_\mu(E) \prod_{i \in \Gamma(E) \cap S} (1 - x_i)^{-1},$$

where x_i's are from (1).

Proof. We prove the inequality by induction on the size of S. The base case is when S is empty and the lemma holds trivially.

For the induction step, let $S_1 = S \cap \Gamma(E)$ and $S_2 = S \setminus S_1$. If $S_1 = \emptyset$, then the lemma holds trivially as E is independent from S in this case. Otherwise S_2 is a proper subset of S. We have that

$$\begin{aligned}
\Pr_\mu(E \mid B(S)) &= \frac{\Pr_\mu(E \wedge B(S_1) \mid B(S_2))}{\Pr_\mu(B(S_1) \mid B(S_2))} \\
&\leq \frac{\Pr_\mu(E \mid B(S_2))}{\Pr_\mu(B(S_1) \mid B(S_2))} \\
&= \frac{\Pr_\mu(E)}{\Pr_\mu(B(S_1) \mid B(S_2))},
\end{aligned}$$

where the last line is because E is independent from $B(S_2)$. We then use the induction hypothesis to bound the denominator. Suppose $S_1 = \{j_1, j_2, \ldots, j_r\}$ for some $r > 0$. Then,

$$\begin{aligned}\Pr_\mu(B(S_1) \mid B(S_2)) &= \Pr_\mu\left(\bigwedge_{i\in S_1} \overline{A_i} \,\middle|\, \bigwedge_{i\in S_2} \overline{A_i}\right)\\ &= \prod_{t=1}^{r} \Pr_\mu\left(\overline{A_{j_t}} \,\middle|\, \bigwedge_{s=1}^{t-1} \overline{A_{j_s}} \wedge \bigwedge_{i\in S_2} \overline{A_i}\right)\\ &= \prod_{t=1}^{r} \left(1 - \Pr_\mu\left(A_{j_t} \,\middle|\, \bigwedge_{s=1}^{t-1} \overline{A_{j_s}} \wedge \bigwedge_{i\in S_2} \overline{A_i}\right)\right).\end{aligned}$$

By the induction hypothesis and (1), we have that for any $1 \le t \le r$,

$$\begin{aligned}\Pr_\mu\left(A_{j_t} \,\middle|\, \bigwedge_{s=1}^{t-1} \overline{A_{j_s}} \wedge \bigwedge_{i\in S_2} \overline{A_i}\right) &\le \Pr_\mu(A_{j_t}) \prod_{i\in\Gamma(j_t)} (1-x_i)^{-1}\\ &\le x_i \prod_{i\in\Gamma(j_t)} (1-x_i) \prod_{i\in\Gamma(j_t)} (1-x_i)^{-1}\\ &= x_i.\end{aligned}$$

Thus,

$$\Pr_\mu(B(S_1) \mid B(S_2)) \ge \prod_{i\in S_1} (1-x_i).$$

The lemma follows. □

Typically we set $x_i = \frac{1}{\Delta+1}$ in the symmetric setting. Then (1) holds if $ep(\Delta+1) \le 1$. In this setting, Lemma 27 is specialized into the following.

Corollary 28. *If $ep(\Delta+1) \le 1$, then*

$$\Pr_\mu(E \mid B(S)) \le \Pr_\mu(E) \left(1 + \frac{1}{\Delta}\right)^{|\Gamma(E)|}.$$

In particular, if $ep(\Delta+1) \le 1$, for any event A_i where $i \notin S$, by Corollary 28,

$$\Pr_\mu(A_i \mid B(S)) \le p_i \left(\frac{\Delta+1}{\Delta}\right)^{\Delta} \le ep. \tag{17}$$

Let Res_t be the resampling set of Algorithm 6 at round $t \ge 1$, and let Bad_t be the set of bad events present at round t. If Algorithm 6 has already stopped at round t, then $\mathsf{Res}_t = \mathsf{Bad}_t = \emptyset$. Furthermore, let $\mathsf{Bad}_0 = \mathsf{Res}_0 = [m]$ since in the first step all random variables are fresh.

Define a random variable $Z_t := C^{-t}\,|\mathsf{Res}_t|$ for $C := 1 - p < 1$ and $t \ge 0$. Note that $Z_0 = C^0\,|[m]| = m$.

Lemma 29. *For any $\Delta \ge 2$, if $6ep\Delta^2 \le 1$ and $3er\Delta \le 1$, then Z_t is a supermartingale.*

Proof. Clearly for any $\Delta \ge 2$, the condition $6ep\Delta^2 \le 1$ implies that $ep(\Delta+1) \le 1$. Therefore the prerequisite of Corollary 28 is met. We will show that Z_t is a supermartingale with respect to the sequence (Res_t) [11, §7.6(3)]. Since $|Z_t| < \infty$, and Z_t is clearly determined by Res_t, we just need to show that

$$\mathbb{E}\left(Z_{t+1} \mid \mathsf{Res}_0, \ldots, \mathsf{Res}_t\right) \le Z_t.$$

Notice that, by Lemma 24, we have that

$$\mathbb{E}\left(Z_{t+1} \mid \mathsf{Res}_0, \ldots, \mathsf{Res}_t\right) = \mathbb{E}\left(Z_{t+1} \mid \mathsf{Res}_t\right).$$

We will show in the following that

$$\mathbb{E}\left(|\mathsf{Res}_{t+1}| \,\middle|\, \text{the set of resampling events at round } t \text{ is (exactly) } \mathsf{Res}_t\right) \le C\,|\mathsf{Res}_t|,$$

where $C = 1 - p$. This implies the lemma.

Let us consider how an event $i \in [m]$ is added into Res_{t+1}. For any $i \in \mathsf{Bad}_{t+1}$, it is always resampled. If $i \in \mathsf{Res}_{t+1} \setminus \mathsf{Bad}_{t+1}$, then it is added due to Algorithm 5 at some round $\ell \ge 1$. We handle the more complicated latter case first.

Call a path $i_0, i_1, \dots, i_\ell$ in the dependency graph G *bad* if the following holds:

(1) $i_0 \in \mathsf{Bad}_{t+1}$;
(2) the event $R_{i_{k-1}i_k}$ holds for every $1 \le k \le \ell$;
(3) any i_k ($k \in [\ell]$) is not adjacent to $i_{k'}$ unless $k' = k-1$ or $k+1$.

Indeed, paths having the third property are induced paths in G. We claim that for any $i \in \mathsf{Res}_{t+1} \setminus \mathsf{Bad}_{t+1}$ added in round $\ell \ge 1$ by Algorithm 5, there exists at least one bad path $i_0, i_1, \dots, i_\ell = i$. We show the claim by an induction on ℓ.

- If $\ell = 1$, then $i = i_1 \in \partial\mathsf{Bad}_{t+1}$ and there must be an $i_0 \in \mathsf{Bad}_{t+1}$ such that (i_0, i_1) is an edge. Since i is marked "resampling" by Algorithm 5, $A_i \cap \sigma_{\mathsf{Bad}_{t+1}} \neq \emptyset$. This implies that $R_{i_0 i_1}$ occurs and the claim holds.
- For the induction step $\ell \ge 2$, due to Algorithm 5, there must exist $i_{\ell-1}$ adjacent to $i_\ell = i$ such that it is marked "resampling" at round $\ell - 1$, and $R_{i_{\ell-1}i_\ell}$ occurs. By the induction hypothesis, there exists a bad path $i_0, \dots, i_{\ell-1}$. Since i is not marked at the round $\ell - 1$, i is not adjacent to any vertices that has been marked up to round $\ell - 2$. Thus i_ℓ is not adjacent to any i_k where $k \le \ell - 2$, and the path $i_0, \dots, i_{\ell-1}, i_\ell$ is bad.

Let $P = i_0, \dots, i_\ell$ be an induced path; that is, for any $k \in [\ell]$, i_k is not adjacent to $i_{k'}$ unless $k' = k-1$ or $k+1$. Only induced paths are potentially bad. Let D_P be the event that P is bad. In other words, $D_P := A_{i_0} \wedge R_{i_0 i_1} \wedge \cdots \wedge R_{i_{\ell-1}i_\ell}$. By Lemma 24, we have that

$$\begin{aligned} &\Pr(P \text{ is bad at round } t+1 \mid \text{the set of resampling events at round } t \text{ is } \mathsf{Res}_t) \\ &= \Pr_\mu(D_P \mid B(\mathsf{Res}_t^e)), \end{aligned} \tag{18}$$

where we recall that we denote $\mathsf{Res}_t^e = [m] \setminus \Gamma^+(\mathsf{Res}_t)$. Applying Corollary 28 with $S = \mathsf{Res}_t^e$, we have that

$$\Pr_\mu(D_P \mid B(\mathsf{Res}_t^e)) \le \Pr_\mu(D_P)\left(1 + \frac{1}{\Delta}\right)^{|\Gamma(D_P)|}. \tag{19}$$

Note that $\Gamma(R_{i_0 i_1}) \subseteq \Gamma^+(A_{i_0})$. By the definition of D_P,

$$\begin{aligned} \Gamma(D_P) &\subseteq \Gamma(A_{i_0}) \cup \Gamma(R_{i_0 i_1}) \cup \cdots \cup \Gamma(R_{i_{\ell-1}i_\ell}) \\ &= \Gamma^+(A_{i_0}) \cup \Gamma(R_{i_1 i_2}) \cup \cdots \cup \Gamma(R_{i_{\ell-1}i_\ell}), \end{aligned}$$

implying that

$$\begin{aligned} |\Gamma(D_P)| &\le \left|\Gamma^+(A_{i_0})\right| + |R_{i_1 i_2}| + \cdots + \left|R_{i_{\ell-1}i_\ell}\right| \\ &\le \ell(\Delta + 1), \end{aligned} \tag{20}$$

as $|\Gamma(R_{ij})| \le |\Gamma^+(A_i)| \le \Delta + 1$ for any $(i, j) \in E$.

We claim that A_{i_0} is independent from $R_{i_{k-1}i_k}$ for any $2 \le k \le \ell$. This is because i_k is not adjacent to i_0 for any $k \ge 2$, implying that

$$\begin{aligned} \mathsf{var}(R_{i_{k-1}i_k}) \cap \mathsf{var}(A_{i_0}) &= \mathsf{var}(A_{i_{k-1}}) \cap \mathsf{var}(A_{i_k}) \cap \mathsf{var}(A_{i_0}) \\ &\subseteq \mathsf{var}(A_{i_k}) \cap \mathsf{var}(A_{i_0}) = \emptyset. \end{aligned}$$

Moreover, any two events $R_{i_{k-1}i_k}$ and $R_{i_{k'-1}i_{k'}}$ are independent of each other as long as $k < k'$. This is also due to the third property of bad paths. Since $k < k'$, we see that $|k' - (k-1)| \ge 2$ and $i_{k'}$ is not adjacent to i_{k-1}. It implies that

$$\begin{aligned} \mathsf{var}(R_{i_{k-1}i_k}) \cap \mathsf{var}(R_{i_{k'-1}i_{k'}}) &= \mathsf{var}(A_{i_{k-1}}) \cap \mathsf{var}(A_{i_k}) \cap \mathsf{var}(A_{i_{k'-1}}) \cap \mathsf{var}(A_{i_{k'}}) \\ &\subseteq \mathsf{var}(A_{i_{k-1}}) \cap \mathsf{var}(A_{i_{k'}}) = \emptyset. \end{aligned}$$

The consequence of these independences is

$$\begin{aligned}\Pr_\mu(D_P) &\le \Pr_\mu(A_{i_0} \wedge R_{i_1 i_2} \wedge \cdots \wedge R_{i_{\ell-1} i_\ell}) \\ &= \Pr_\mu(A_{i_0}) \prod_{k=2}^{\ell} \Pr_\mu(R_{i_{k-1} i_k}) \\ &\le p r^{\ell-1}. \end{aligned} \tag{21}$$

Note that in the calculation above we ignore $R_{i_0 i_1}$ as it can be positively correlated to A_{i_0}.

Combining (18), (19), (20), and (21), we have that

$$\begin{aligned}&\Pr(D_P \mid \text{the set of resampling events at round } t \text{ is (exactly) } \mathsf{Res}_t) \\ &\le p r^{\ell-1} \left(1+\frac{1}{\Delta}\right)^{\ell(\Delta+1)} \le \frac{p}{r}\left(\left(1+\frac{1}{\Delta}\right) er\right)^{\ell}. \end{aligned} \tag{22}$$

In order to apply a union bound on all bad paths, we need to bound their number. The first vertex i_0 must be in Bad_{t+1}, implying that $i_0 \in \Gamma^+(\mathsf{Res}_t)$. Hence there are at most $(\Delta+1)\,|\mathsf{Res}_t|$ choices. Then there are at most Δ choices of i_1 and $(\Delta-1)$ choices of every subsequent i_k where $k \ge 2$. Hence, there are at most $\Delta(\Delta-1)^{\ell-1}$ induced paths of length $\ell \ge 1$, originating from a particular $i_0 \in \Gamma^+(\mathsf{Res}_t)$. Thus, by a union bound on all potentially bad paths and (22),

$$\begin{aligned}&\mathbb{E}\left(\,|\mathsf{Res}_{t+1} \setminus \mathsf{Bad}_{t+1}| \;\middle|\; \text{the set of resampling events at round } t \text{ is (exactly) } \mathsf{Res}_t\right) \\ &\le \sum_{\ell=1}^{\infty} (\Delta+1)\,|\mathsf{Res}_t|\,\Delta(\Delta-1)^{\ell-1} p/r \left(\left(1+\frac{1}{\Delta}\right) er\right)^{\ell} \\ &= \frac{(\Delta+1)\Delta p}{(\Delta-1)r}\,|\mathsf{Res}_t| \sum_{\ell=1}^{\infty} \left(\left(\frac{\Delta^2-1}{\Delta}\right) er\right)^{\ell} \\ &\le \frac{(\Delta+1)\Delta p}{(\Delta-1)r}\,|\mathsf{Res}_t| \sum_{\ell=1}^{\infty} (er\Delta)^{\ell} = \frac{(\Delta+1)\Delta p}{(\Delta-1)r} \cdot \frac{er\Delta}{1-er\Delta}\,|\mathsf{Res}_t| \\ &= \frac{\Delta+1}{\Delta-1} \cdot \frac{3}{2} \cdot ep\Delta^2\,|\mathsf{Res}_t|\,, \end{aligned} \tag{23}$$

where we use the condition that $er\Delta \le 1/3$.

On the other hand, it is straightforward to bound the size of $\mathsf{Bad}_{t+1} \subseteq \Gamma^+(\mathsf{Res}_t)$. If $i \in \mathsf{Bad}_{t+1}$, then there are two possibilities. The first scenario is that $i \in \mathsf{Res}_t$ and then all of its random variables are fresh. In this case it occurs with probability $p_i \le p$. Otherwise $i \in \partial\mathsf{Res}_t$. Recall that by Lemma 24, the distribution at round $t+1$ is $\Pr_\mu(\cdot \mid B(\mathsf{Res}_t^e))$. By Corollary 28, for any $i \in \partial\mathsf{Res}_t$,

$$\Pr_\mu\left(A_i \mid B(\mathsf{Res}_t^e)\right) \le p\left(1+\frac{1}{\Delta}\right)^{\Delta} \le ep.$$

This implies that

$$\begin{aligned}&\mathbb{E}\left(|\mathsf{Bad}_{t+1}| \mid \text{the set of resampling events at round } t \text{ is (exactly) } \mathsf{Res}_t\right) \\ &\le p\,|\mathsf{Res}_t| + ep\,|\partial\mathsf{Res}_t| \le p(1+e\Delta)\,|\mathsf{Res}_t|\,. \end{aligned} \tag{24}$$

Combining (23) and (24), we have that

$$\begin{aligned}&\mathbb{E}\left(|\mathsf{Res}_{t+1}| \mid \text{the set of resampling events at round } t \text{ is (exactly) } \mathsf{Res}_t\right) \\ &\le \frac{\Delta+1}{\Delta-1} \cdot \frac{3}{2} \cdot ep\Delta^2\,|\mathsf{Res}_t| + p(1+e\Delta)\,|\mathsf{Res}_t| \\ &= p\left(\frac{\Delta+1}{\Delta-1} \cdot \frac{3}{2} \cdot e\Delta^2 + (1+e\Delta)\right)|\mathsf{Res}_t|\,. \end{aligned}$$

All that is left is to verify that

$$p\left(\frac{\Delta+1}{\Delta-1}\cdot\frac{3}{2}\cdot e\Delta^2+(1+e\Delta)\right)\leq C.$$

This is straightforward by the condition $6ep\Delta^2\leq 1$ and $\Delta\geq 2$, as

$$\begin{aligned}C-p\left(\frac{\Delta+1}{\Delta-1}\cdot\frac{3}{2}\cdot e\Delta^2+(1+e\Delta)\right)&\geq 6ep\Delta^2-p-p\left(\frac{\Delta+1}{\Delta-1}\cdot\frac{3}{2}\cdot e\Delta^2+(1+e\Delta)\right)\\&\geq p\left(6e\Delta^2-1-\frac{\Delta+1}{\Delta-1}\cdot\frac{3}{2}\cdot e\Delta^2-(1+e\Delta)\right)\geq 0.\quad\square\end{aligned}$$

Since Z_t is a supermartingale by Lemma 29, we have, by the tower property of expectations [11, §7.7(16)],

$$\mathbb{E}\,Z_t\leq\mathbb{E}\,Z_0=m.$$

In other words, $\mathbb{E}\,|\mathsf{Res}_t|\leq C^t m$. As $C<1$, the expected number of resampling events is

$$\sum_{t=0}^{\infty}\mathbb{E}\,|\mathsf{Res}_t|\leq\sum_{t=0}^{\infty}C^t m=\frac{1}{1-C}\cdot m.$$

This implies the first part of Theorem 26. For the second part, just observe that within $O(\log m)$ rounds, the expected number of bad events is less than 1.

The first condition of Theorem 26 requires p to be roughly $O(\Delta^{-2})$. This is necessary, due to the hardness result in [4] (see also Theorem 33). Also, in the analysis, it is possible to always add all of $\partial\mathsf{Bad}_t$ into Res_t. Consider a monotone CNF formula. If a clause is unsatisfied, then all of its neighbours need to be added into the resampling set. Such behaviours would eventually lead to the $O(\Delta^{-2})$ bound. This situation is in contrast to the resampling algorithm of Moser and Tardos [20], which only requires $p=O(\Delta^{-1})$ as in the symmetric Lovász Local Lemma.

7. Applications of Algorithm 6

7.1. **k-CNF Formulas.** Consider a k-CNF formula where every variable appears in at most d clauses. Then Theorem 1 says that if $d\leq 2^k/(ek)+1$, then there exists a satisfying assignment. However, as mentioned in Section 4.3, [4, Corollary 30] showed that when $d\geq 5\cdot 2^{k/2}$, then sampling satisfying assignments is **NP**-hard, even restricted to monotone formulas.

To apply Algorithm 6 in this setting, we need to bound the parameter r in Theorem 26. A natural way is to lower bound the number of shared variables between any two *dependent* clauses. If this lower bound is s, then $r=2^{-s}$ since there is a unique assignment on these s variables that can be extended in such a way as to falsify the clauses.

Definition 30. *Let $d\geq 2$ and $s\geq 1$. A k-CNF formula is said to have* degree d *if every variable appears in at most d clauses. Moreover, it has* intersection s *if for any two clauses C_i and C_j that share at least one variable, $|\mathsf{var}(C_i)\cap\mathsf{var}(C_j)|\geq s$.*

Note that by the definition if $k<s$ then the formula is simply isolated clauses. Otherwise, $k\geq s$ and we have that $p_i=p=2^{-k}$ and $r\leq 2^{-s}$. The maximum degree Δ in the dependency graph is at most $(d-1)k$. Recall Theorem 26. The condition $6ep\Delta^2\leq 1$ becomes

$$d\leq\frac{1}{\sqrt{6ek}}\cdot 2^{k/2}+1,$$

and $3er\Delta\leq 1$ becomes

$$s\geq\log(d-1)+\log k+\log(3e).$$

Thus by Theorem 26 we have the following result. Note that each resampling event involves at most k variables, and for k-CNF formulas with degree d, the number of clauses is linear in the number of variables.

Corollary 31. *If* $d \le \frac{1}{\sqrt{6}ek} \cdot 2^{k/2} + 1$ *and* $s \ge \log(d-1) + \log k + \log(3e)$*, then Algorithm 6 samples satisfying assignments of* k*-CNF formulas with degree* d *and intersection* s *in expected* $O(n)$ *time where* n *is the number of variables.*

We also show that restricting the number of shared variables does not make the problem trivial. Indeed, the bad instance in the proof of [4, Corollary 30] has roughly $k/2$ shared variables for each pair of dependent clauses. For completeness, we will show that if k is even, and $d \ge 4 \cdot 2^{k/2}$ and $s = k/2$, then the sampling problem is **NP**-hard. The proof is almost identical to that of [4, Corollary 30]. The case of odd k can be similarly handled but with larger constants.

We will use the inapproximability result of Sly and Sun [25] (or equivalently, of Galanis et al. [10]) for the hard-core model. We first remind the reader of the relevant definitions. Let $\lambda > 0$. For a graph $G = (V, E)$, the hard-core model with parameter $\lambda > 0$ is a probability distribution over the set of independent sets of G; each independent set I of G has weight proportional to $\lambda^{|I|}$. The normalizing factor of this distribution is the partition function $Z_G(\lambda)$, formally defined as $Z_G(\lambda) := \sum_I \lambda^{|I|}$ where the sum ranges over all independent sets I of G. The hardness result we are going to use is about approximating $Z_G(\lambda)$, but it is standard to translate it into the sampling setting as the problem is self-reducible.

Theorem 32 ([25, 10])**.** *For* $d \ge 3$*, let* $\lambda_c(d) := (d-1)^{d-1}/(d-2)^d$*. For all* $\lambda > \lambda_c(d)$*, it is* **NP***-hard to sample an independent set* I *with probability proportional to* $\lambda^{|I|}$ *in a* d*-regular graph.*

Theorem 33. *Let* k *be an even integer. If* $d \ge 4 \cdot 2^{k/2}$ *and* $s = k/2$*, then it is* **NP***-hard to sample satisfying assignments of* k*-CNF formulas with degree* d *and intersection* s *uniformly at random.*

Proof. Given a d-regular graph $G = (V, E)$, we will construct a monotone k-CNF formula C with degree d and intersection $k/2$ such that satisfying assignments of C can be mapped to independent sets of G. Replace each vertex $v \in V$ by s variables, say $v_1, \dots, v_s$. If $(u, v) \in E$, then create a monotone clause $v_1 \vee \cdots \vee v_s \vee u_1 \vee \cdots \vee u_s$. It is easy to see that every variable appears exactly d times since G is d-regular. Moreover, the number of shared variables is always s and the clause size is $2s = k$.

For each satisfying assignment, we map it to a subset of vertices of G. If all of $v_1, \dots, v_s$ are *false*, then make v occupied. Otherwise v is unoccupied. Thus a satisfying assignment is mapped to an independent set of G. Moreover, there are $(2^{k/2} - 1)^{n-|I|}$ satisfying assignments corresponding to an independent set I, where n is the number of vertices in G. Thus the weight of I is proportional to $(2^{k/2} - 1)^{-|I|}$; namely $\lambda = (2^{k/2} - 1)^{-1}$ in the hard-core model.

In order to apply Theorem 32, all we need to do is to verify that $\lambda > \lambda_c$, or equivalently

$$2^{k/2} - 1 < \frac{(d-2)^d}{(d-1)^{d-1}}.$$

This can be done as follows,

$$\frac{(d-2)^d}{(d-1)^{d-1}} = (d-2)\Big(1 - \frac{1}{d-1}\Big)^{d-1} \ge \Big(\frac{4}{5}\Big)^5 (d-2) > 2^{k/2} - 1. \qquad \square$$

Due to Theorem 33, we see that the dependence between k and d in Corollary 31 is tight in the exponent, even with the further assumption on intersection s.

7.2. **Independent Sets.** We may also apply Algorithm 6 to sample hard-core configurations with parameter λ. Every vertex is associated with a random variable which is occupied with probability $\frac{\lambda}{1+\lambda}$. In this case, each edge defines a bad event which holds if both endpoints are occupied. Thus $p = \left(\frac{\lambda}{1+\lambda}\right)^2$. Algorithm 6 is specialized to Algorithm 7.

To see this, consider a graph $G = (V, E)$ with maximum degree d. Given a configuration $\sigma : V \to \{0, 1\}$, consider the subgraph $G[\sigma]$ of G induced by the vertex subset $\{v \in V : \sigma(v) = 1\}$.

Algorithm 7 Sample Hard-core Configurations

(1) Mark each vertex occupied with probability $\frac{\lambda}{1+\lambda}$ independently.

(2) While there is at least one edge with both end points occupied, resample all occupied components of sizes at least 2 and their boundaries.

(3) Output the set of vertices.

Then we denote by $\mathsf{BadVtx}(\sigma)$ the set of vertices in any component of $G[\sigma]$ of size at least 2. Then the output of Algorithm 5 is

$$\mathsf{ResVtx}(\sigma) := \mathsf{BadVtx}(\sigma) \cup \partial\mathsf{BadVtx}(\sigma).$$

This is because first, all of $\partial\mathsf{BadVtx}(\sigma)$ will be resampled, since any of them has at least one occupied neighbour in $\mathsf{BadVtx}(\sigma)$. Secondly, $v \in \partial\mathsf{BadVtx}(\sigma)$ is unoccupied (otherwise $v \in \mathsf{BadVtx}(\sigma)$), and Algorithm 5 stops after adding all of $\partial\mathsf{BadVtx}(\sigma)$. This explains Algorithm 7.

Moreover, let $\mathsf{Bad}(\sigma)$ be the set of edges whose both endpoints are occupied under σ. Let $\mathsf{Res}(\sigma)$ be the set of edges whose both endpoints are in $\mathsf{ResVtx}(\sigma)$. Let σ_t be the random configuration of Algorithm 7 at round t if it has not halted, and $\mathsf{Bad}_t = \mathsf{Bad}(\sigma_t)$, $\mathsf{Res}_t = \mathsf{Res}(\sigma_t)$.

Lemma 34. *If $ep(2d-1) < 1$, then $\mathbb{E}\,|\mathsf{Bad}_{t+1}| \le (4ed^2 - 1)p\,\mathbb{E}\,|\mathsf{Bad}_t|$.*

Proof. First note that the dependency graph is the line graph of G and $\Delta = 2d-2$ is the maximum degree of the line graph of G. Thus $ep(2d-1) < 1$ guarantees the prerequisite of Corollary 28 is met. It also implies that for any σ, $|\mathsf{Res}(\sigma)| \le (2\Delta - 1)\mathsf{Bad}(\sigma)$, and $\partial\,|\mathsf{Res}(\sigma)| \le (2d-2)\,|\mathsf{Res}(\sigma)|$. Similarly to the analysis in Lemma 29, conditioned on a fixed Bad_t, by Corollary 28 (or (17) in particular), we have that

$$\begin{aligned}\mathbb{E}\,|\mathsf{Bad}_{t+1}| &\le p\,|\mathsf{Res}(\sigma)| + ep\,|\partial\mathsf{Res}(\sigma)| \\ &\le (p(2d-1) + ep(2d-2)(2d-1))\,|\mathsf{Bad}_t| \\ &< (4ed^2 - 1)p\,|\mathsf{Bad}_t|\,.\end{aligned}$$

Since the inequality above holds for any Bad_t, the lemma follows. □

Lemma 34 implies that, if $4epd^2 \le 1$, then the number of bad edges shrinks with a constant factor, and Algorithm 7 resamples $O(m)$ edges in expectation where $m = |E|$. A bounded degree graph is sparse and thus the expected running time of Algorithm 7 is $O(n)$. Since $p = \left(\frac{\lambda}{1+\lambda}\right)^2$, the condition $4epd^2 \le 1$ is equivalent to

$$\lambda \le \frac{1}{2\sqrt{e}d - 1}.$$

Thus we have the following theorem, where the constants are slightly better than directly applying Theorem 26.

Theorem 35. *If $\lambda \le \frac{1}{2\sqrt{e}d-1}$, then Algorithm 7 draws a uniform hard-core configuration with parameter λ from a graph with maximum degree d in expected $O(n)$ time.*

The optimal bound of sampling hard-core configurations is $\lambda < \lambda_c \approx \frac{e}{d}$ where λ_c is defined in Theorem 32. The algorithm is due to Weitz [27] and the hardness is shown in [25, 10]. The condition of our Theorem 35 is more restricted than correlation decay based algorithms [27] or traditional Markov chain based algorithms. Nevertheless, our algorithm matches the correct order of magnitude $\lambda = O(d^{-1})$. Moreover, our algorithm has the advantage of being simple, exact, and running in linear time in expectation.

References

[1] Dimitris Achlioptas and Fotis Iliopoulos. Random walks that find perfect objects and the Lovász local lemma. *J. ACM*, 63(3):22, 2016. 3

[2] Noga Alon. A parallel algorithmic version of the local lemma. *Random Struct. Algorithms*, 2(4):367–378, 1991. 1
[3] József Beck. An algorithmic approach to the Lovász local lemma. I. *Random Struct. Algorithms*, 2(4):343–366, 1991. 1
[4] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, Heng Guo, and Daniel Štefankovič. Approximation via correlation decay when strong spatial mixing fails. In *ICALP,* 45:1–13, 2016. 1, 2, 14, 17, 18, 22, 23
[5] Magnus Bordewich, Martin E. Dyer, and Marek Karpinski. Stopping times, metrics and approximate counting. In *ICALP*, pages 108–119, 2006. 3
[6] Russ Bubley and Martin E. Dyer. Graph orientations with no sink and an approximation for a hard case of #SAT. In *SODA*, pages 248–257, 1997. 11
[7] Henry Cohn, Robin Pemantle, and James G. Propp. Generating a random sink-free orientation in quadratic time. *Electr. J. Comb.*, 9(1), 2002. 1, 4, 11
[8] Artur Czumaj and Christian Scheideler. Coloring nonuniform hypergraphs: A new algorithmic approach to the general Lovász local lemma. *Random Struct. Algorithms*, 17(3-4):213–237, 2000. 1
[9] Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Infinite and finite sets, volume 10 of Colloquia Mathematica Societatis János Bolyai*, pages 609–628, 1975. 1, 3
[10] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability & Computing*, 25(4):500–559, 2016. 3, 23, 24
[11] Geoffrey R. Grimmett and David R. Stirzaker. *Probability and random processes*. Oxford University Press, New York, third edition, 2001. 19, 22
[12] David G. Harris and Aravind Srinivasan. A constructive algorithm for the Lovász local lemma on permutations. In *SODA*, pages 907–925, 2014. 3
[13] Nicholas J. A. Harvey and Jan Vondrák. An algorithmic proof of the Lovász local lemma via resampling oracles. In *FOCS*, pages 1327–1346, 2015. Full version available at abs/1504.02044. 3, 9
[14] Jonathan Hermon, Allan Sly, and Yumeng Zhang. Rapid mixing of hypergraph independent set. *CoRR*, abs/1610.07999, 2016. 3
[15] Donald E. Knuth. The art of computer programming, volume 4b (draft, pre-fascicle 6a). 2015. Available at `http://www-cs-faculty.stanford.edu/~uno/fasc6a.ps.gz`. 9
[16] Kashyap Babu Rao Kolipaka and Mario Szegedy. Moser and Tardos meet Lovász. In *STOC*, pages 235–244, 2011. 2, 4, 9
[17] Jingcheng Liu and Pinyan Lu. FPTAS for counting monotone CNF. In *SODA*, pages 1531–1548, 2015. 3
[18] Ankur Moitra. Approximate counting, the Lovász local lemma and inference in graphical models. *CoRR*, abs/1610.04317, 2016. 3
[19] Michael Molloy and Bruce A. Reed. Further algorithmic aspects of the local lemma. In *STOC*, pages 524–529, 1998. 1
[20] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2), 2010. 1, 4, 5, 22
[21] James G. Propp and David B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algorithms*, 9(1-2):223–252, 1996. 2
[22] James G. Propp and David B. Wilson. How to get a perfectly random sample from a generic markov chain and generate a random spanning tree of a directed graph. *J. Algorithms*, 27(2):170–217, 1998. 13
[23] Alexander D. Scott and Alan D. Sokal. The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *Journal of Statistical Physics*, 118(5):1151–1261, 2005. 10
[24] James B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985. 2, 3
[25] Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d-regular graphs. *The Annals of Probability*, 42(6):2383–2416, 2014. 3, 23, 24
[26] Aravind Srinivasan. Improved algorithmic versions of the Lovász local lemma. In *SODA*, pages 611–620, 2008. 1
[27] Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149, 2006. 3, 24
[28] David B. Wilson. Generating random spanning trees more quickly than the cover time. In *STOC*, pages 296–303, 1996. 1, 4, 5, 13

School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom.
E-mail address: `h.guo@qmul.ac.uk`

School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom.
E-mail address: `mj@qmul.ac.uk`

Department of EECS, University of California, Berkeley, CA
E-mail address: `liuexp@berkeley.edu`